The rise and fall of the middle class:

technology, skills, and inequality

Luis Valenzuela R.

Wolfson College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2016

Acknowledgements

First and foremost I am grateful to God for all the blessings I have received, which have enabled me to embark successfully into a project like this one. I would like to thank my supervisor, David Vines, for his constant and patient guidance over all the years this project was built on. David was able to adapt my rather chaotic research approach into a coherent project. Working alongside him has been both enjoyable and rewarding; David is not only a fantastic academic, but a great mentor. I thank Martin Ellison for his insightful comments, which gave rise to important part of the work I include here. I thank particularly my friends and colleagues Michael Koelle and Manuel Fernandez, from whose many comments my work benefited. I would like to thank my family, for their constant support over the years, which I have felt close despite the long distance. I also thank my new family here in England, who have made me feel much like at home. Finally, I am grateful to Anna, who accompanied me all along this process, always with warm advice, care and love.

Abstract

Title: The rise and fall of the middle class: technology, skills, and inequality

By: Luis Valenzuela R., Wolfson College, University of Oxford

For: Doctor of Philosophy, Michaelmas 2016

Over the twentieth century advanced economies have seen an economic and social development process which was build upon the consolidation of a strong middle class. Yet, recent decades have seen an increase in wealth and income inequality, reaching levels not seen since before the Second World War. This thesis explore some of the these issues in two parts. The focus of the first part of the thesis is on the role of education and technology in the *rise* of the middle class. By means of an overlapping generations model with endogenous growth, I study the conditions that enable a society to transit from underdevelopment to development. The model in place reproduces a Kuznets curve, which is deemed an important empirical feature of the history of advanced economies. The second part focuses on the *fall* of the middle class, by studying the effect of technology and skills in job polarisation – i.e. the fall in employment in middle-skill occupations. The approach is both theoretical and empirical. A sorting model based on *tasks* is developed and adapted to study polarisation. Central to this model are the distributions of skills that workers have. Thus, a complete chapter is dedicated to characterise these ability distributions, using longitudinal data from the UK for 1991-2008 in an econometric model based on the so-called Mincer equation. The estimated distributions – positive skewed – are used to calibrate the sorting model. Then, this model is used to identify the nature of the technological process affecting the UK economy over the selected period of study. Simple counterfactual exercises shed light on the strong effect of technical progress on both polarisation and inequality. In contrast, the role of change in skills is negligible. The overall conclusion is that the nature of technological change is essential in defining distributional outcomes: whilst technology can enable the rise to a strong middle class, it can also undermine it.

Word count (including appendix, excluding bibliography): $\approx 72{,}000$ (counted using TeXcount software)

Contents

List of Figures

Introduction

"Thus it is manifest that the best political community is formed by citizens of the middle class, and that those states are likely to be well-administered in which the middle class is large, and stronger if possible than both the other classes, or at any rate than either singly; for the addition of the middle class turns the scale, and prevents either of the extremes from being dominant."

– Aristotle, *Politics (Book Four, Part XI)*

0.1 The historical rise of the middle class

Living standards have been steadily rising in most of (now) advanced economies, at least since the outbreak of the industrial revolution (Maddison, 2003). This increase was particularly rapid after the Second World War, as Figure 0.1 shows. Together with these changes in GDP, there has been a consequent increase in consumption, life expectancy, levels of education, and many other social indicators.[1]

Alongside this aggregate phenomenon, it is interesting to understand how generalised this increase in welfare has been. The simple categorisation used here is that of the "middle class", the rich, and the poor.[2] This distinction is rather loose and very pragmatic. For example, the "middle class" can be defined as a fixed proportion of the total population (e.g. the middle 80%, as in Easterly, 2001), or as the variable proportion of those earning

[1]For example, see Roser (2016), and the related comprehensive database on economic and social indicators at https://ourworldindata.org/

[2]There is considerable disagreement within the economics literature and between economics and sociology about the definition of the middle class. The approach here is very pragmatic. For a recent discussion on concepts and measurements in economics, see Jayadev, Lahoti and Reddy (2015).

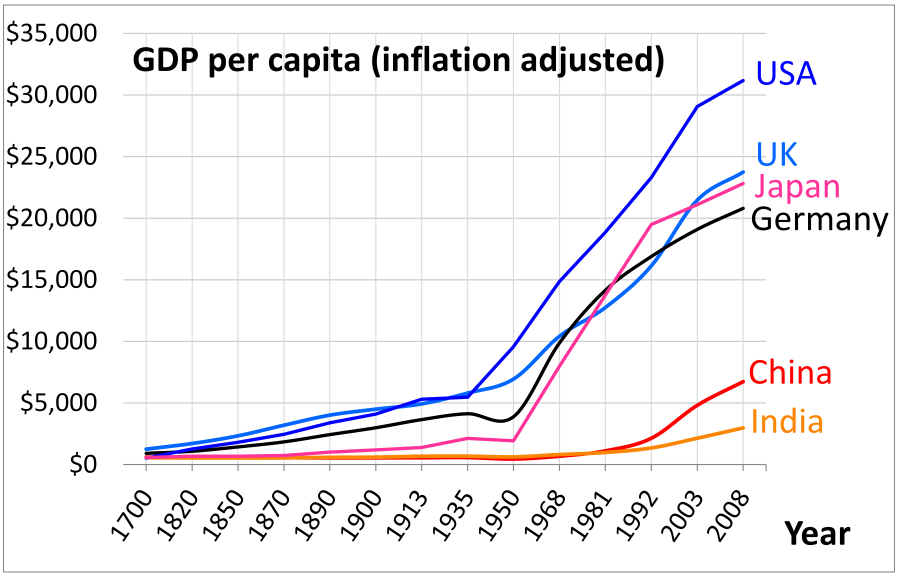

Figure 0.1: Historical evolution of GDP per capita in selected countries, from Maddison (2003)

between a certain income interval (as in Banerjee and Duflo, 2008). Whatever the case, our interest is with the evolution of the income, wealth or welfare of the middle class, which the fixed-proportion definition captures directly, and the variable-proportion definition captures indirectly, through changes in the size of this group.

Council of Economic Advisers (2015) provides an interest summary of the evolution of several variables for the bottom 90% of the population, for a selection of countries.[3] Figure 0.2 presents one of their findings. Namely, there was a remarkable increase in real income for the bottom 90% of the population of selected countries for at least three decades after the Second World War (notice the log scale of the graph). They also report a rapid annual growth in labour productivity for this group, for the same countries and period.

The data reported before only starts in 1950. Analysis for earlier decades is usually available from studies focusing on top income shares. For example, Alvaredo et al. (2013) describe the evolution of the income share of the Top 1 percent of the population (for whom tax-returns are available). The main result is shown in Figure 0.3, equivalent to Figure 2A in the original paper. There is a clear fall in the income of the highest

[3]Clearly, the bottom 90% is not equivalent to any definition of the middle class, as they also include the poorest in the economy. Still, from a pure statistically point of view, the income of the bottom 90% with and without the bottom 10% is likely to be very similar.

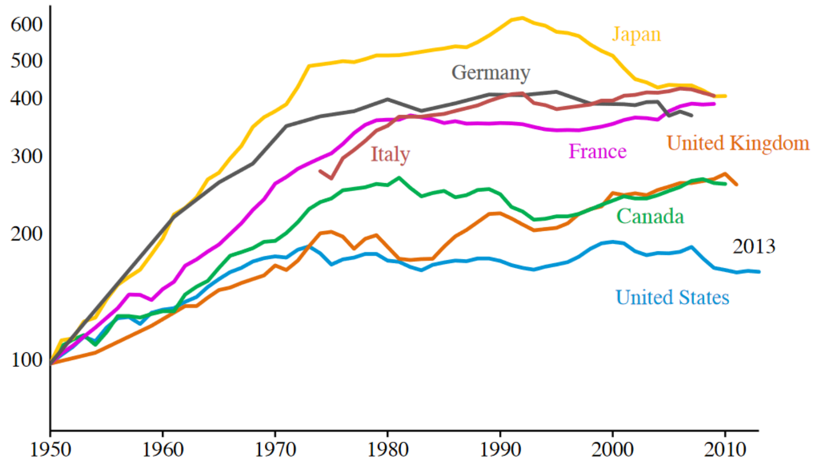

Figure 0.2: Growth in Real Average Income for the Bottom 90%, 1950–2013, from Council of Economic Advisers (2015) (log scale, 1950=100)

earning individuals at least until 1980, for the four selected countries. The story is roughly the same for the top 10%. Since GDP per capita was growing fast over the period, and the income share of the richest agents was decreasing, necessarily the share of income of the middle and bottom was going up. This is consistent also with the analysis of the bottom of the distribution. Bourguignon and Morrisson (2002) shows how both poverty and extreme poverty has fallen consistently since 1820 (when their data begin) until our days.

The evolution of inequality provides further evidence of the widespread benefits of economic growth. Kuznets (1955, 1963) documents a decline in income inequality between the late nineteenth century (UK) or early twentieth century (US) until the post-war period. Kuznets attempts to provide an explanation of this trends, which came to be known as the Kuznets curve (an inverse U-shaped relationship between inequality and GDP). With much better data, Milanovic (2016) is among the latest to show a clear Kuznets curve in both UK and US between at least the 17th century, with the peak on inequality around 1860, and minimum around 1970s. This fall in income inequality throughout most of the twentieth century (with a short-lived spike during the Great Depression) while GDP per capita was rapidly increasing, suggests an improvement in incomes for a great portion of the population.

The evidence presented so far, albeit indirect, supports the view that an important part of the aggregate improvement in living standards as measured from GDP has accrued

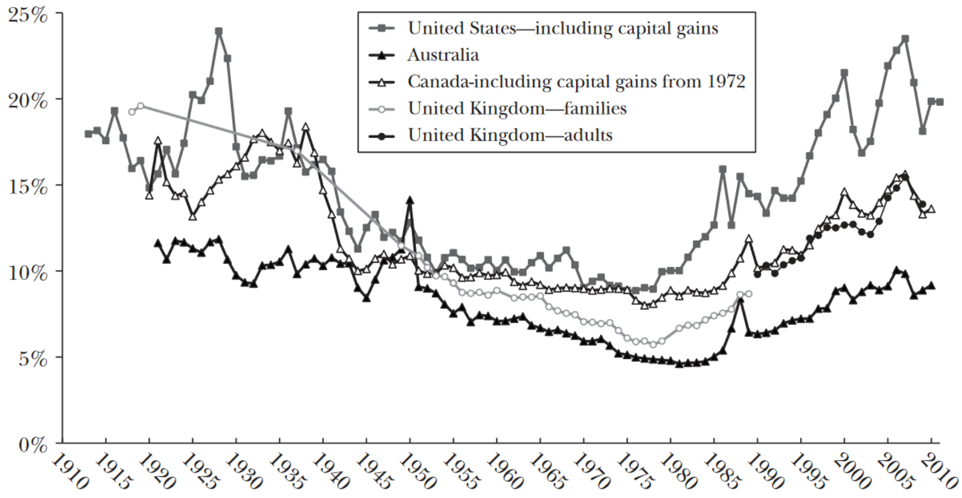

Figure 0.3: Top 1% income shares in selected countries, from Alvaredo et al. (2013), reproduced with permission of the American Economic Association

to the bottom and middle classes of the population (if defined as a fixed proportion of the population); equivalently, this increase in income for the many meant an increase in the size of the middle class (if defined as a variable proportion of the population).

There is no agreement in the literature about the drivers of the developments just documented. Many theories have been proposed, which this introduction does not attempt to review. The rapid increase in living standards can be seen to be an effect of strong technological change, which has increased the productivity of capital and labour (e.g. Kremer, 1993; Maddison, 2003). Additionally, education and the accumulation of human capital have been remarkable over the period. For example, Lee and Lee (2016) calculate that the primary education enrolment rate in advanced economies grew from 30% in 1820 to almost 90% by 1940, including both male and female. Secondary education enrolment rates grew quite fast too, albeit only since the beginning of the twentieth century. These authors also set out to calculate the human capital stock per worker for advanced (and developing) economies since 1870 (Figure 11 in original paper). This figure reveals a continuing increase since then. This brings to mind the literature on endogenous growth that Paul Roemer and Robert Lucas have pioneered. Their theories highlight the positive effect of the accumulation of knowledge on the productivity of factors of production – which have decreasing marginal returns. Last but not least, the large and relatively new literature on Institutions is also prominent in putting forward the crucial role which pro-growth institutions have played in the development process, including via capital accumulation, FDI, trade openness, and a number of other different ways (e.g. Acemoglu, Johnson and Robinson, 2004).

There is yet one factor which played a particularly important role in the transmission of increases in income towards the middle of the distribution. This is the taxation system. Alvaredo et al. (2013) show that the top marginal income tax rates increased sharply from 0% around 1900s to as high as 100% (US) in 1940, remaining high until the 1980s, in particular the US and UK. On the one hand, for a given level of government revenues, higher taxation of top earners reduces the need to tax less wealthy individuals, leading to more progressive tax structures. On the other hand, high tax revenues increases the resources available to fund public mass education efforts and other public services that benefit a greater number of people, contributing towards the creation of the now called Welfare State. Certainly, the aforementioned expansion in GDP and the associated increase in the taxation base also enhances the capacity of the state to provide for its less well-off citizens, further contributing to the "rise" of the middle class. All of the above is partly a consequence of the increased political representation of lower classes in the twentieth century, reflected in a legislation more tilted toward their interests.

0.2 The more recent fall of the middle class

The aforementioned expansion in the middle class has seen recent setbacks, in particular since the 1980s. The same evidence presented above confirms this. In particular, the growth of income accruing to the bottom 90% has if anything halted since 1980, even when GDP per capita has continued to increase (see Figures 0.1 and Figure 0.2). Additionally, Council of Economic Advisers (2015) calculates a consistent slowdown in labour productivity growth for the bottom 90% at least since 1965. The recent decoupling between average earnings and productivity growth is also well documented (Sherk, 2013). All this is consistent with a greater share of total income being captured by the top 10% and 1%, as Figure 0.3 indicates.

Another well-documented development affecting advanced economies and the middle class is the observed shift in employment from mid-skill intensive jobs to **both** low

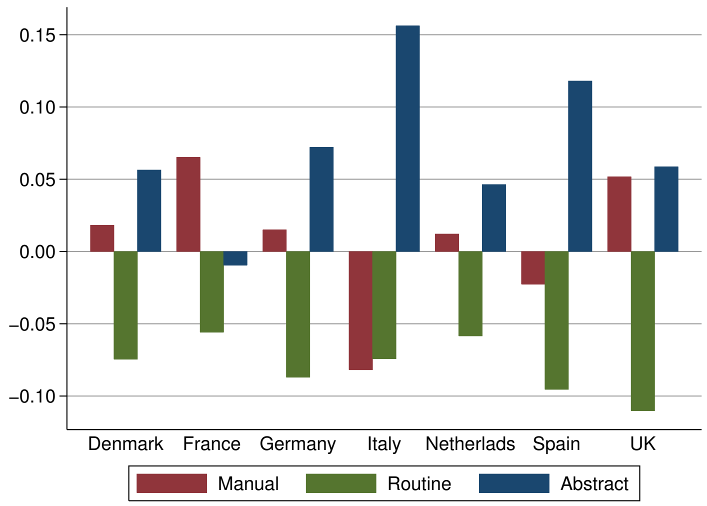

Figure 0.4: Change in employment shares, 1992-2008, selected EU countries

and high skill intensive jobs, at least since the 1980s (Autor, Katz and Kearney, 2006; Acemoglu and Autor, 2011; Goos, Manning and Salomons, 2014). This hollowing out of the labour markets is known in the literature as job polarisation. In Figure 0.4 I present further evidence of this, for some of the largest EU countries, using harmonised employment data from Eurostat. In particular it shows the change in employment shares for three occupational groups, between 1992 and 2008. The categorisation of these groups is based on a conceptualisation of the "tasks" that workers can perform in every job. These tasks are "manual", "routine", and "abstract", as common in the literature (e.g. Autor, 2013). Every occupation in the economy is mapped to one of these tasks, thereby classifying all workers in the economy.[4] This three tasks' categorisation is expected to reflect low-skill ("manual"), mid-skill ("routine") and high-skill ("abstract") workers. The case of polarisation is clear in several countries like Denmark, Germany, the Netherlands, and particularly the UK. In other countries there is also a considerable fall in mid-skill employment, but the direction of reallocation is mainly in one direction, either up (Italy or Spain) or down (France).

Atkinson and Brandolini (2011) provide a comprehensive analysis of the changes facing the middle class in advanced economies in the recent decades. Figure 0.5 reproduces

[4] The concept ot tasks and the three categories used here are studied in Chapter 2. Notice that the Eurostat only provides information on occupations at the 1-digit level, in direct contrast with the datasets used in Chapter 3. The mapping between tasks and occupations is then rather crude, with the majority of the problems arising from assuming some routine jobs as manual.

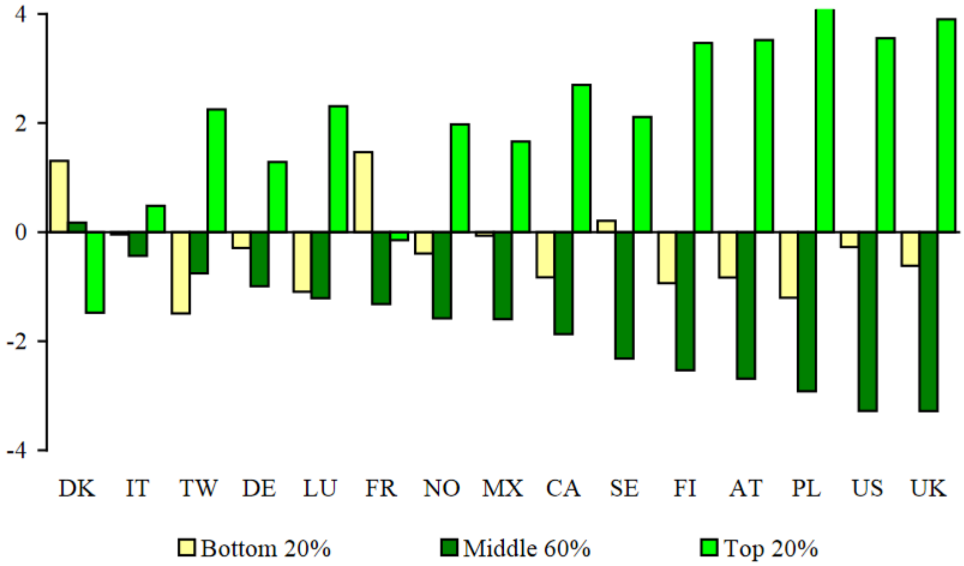

Source: Author's calculations from the LIS database, as of 10 May 2011.

Figure 0.5: Percentage change in income shares for three groups, 1985-2004, selected countries, from Atkinson and Brandolini (2011), reproduced with authors' permission

Figure 1 in their paper, showing the change in income shares of three groups, for selected countries.[5] In particular, they use a fixed proportion definition of middle class, defined as the middle 60% of the population. In the majority of countries but Denmark, the income share of the middle has fallen dramatically. In most of the cases the income of the bottom 20% has also fallen, with a corresponding large rise in the proportion of total income accruing to the top 20%. Figure 5 in Atkinson and Brandolini (2011) (not shown here) also documents the strong positive correlation between an index of polarisation and the Gini index over the same period, thereby providing an interesting link between the two phenomena.

Alongside this "fall" in the middle class, there has been a surge in inequality, as has been widely documented in the literature (see Atkinson and Morelli, 2014 for a data-intensive summary). Piketty (2014) is perhaps a remarkable example, also documenting a considerable rise in income and wealth inequality in many advanced economies since the 1980s. More recently, Milanovic (2016) suggested the existence of Kuznets *waves*, of which – he claims – we are currently living the upward phase.

[5]The data comes from the Luxembourg Income Study Database (LIS). The country codes correspond to those in the LIS database, which meaning can be found at http://www.lisdatacenter.org/our-data/lis-database/documentation/list-of-datasets/

The factors driving this erosion of the middle class are varied. From a more broad perspective, there are important political factors at play, with the rise of neoliberalism in countries like the US and UK, which has lowered taxation rates for top incomes (Alvaredo et al., 2013), reduced the power of trade unions, and deregulated labour markets. Furthermore, CEO pay has grown much faster than median incomes in many advanced economies at least since the 1980s. Frydman and Jenter (2010) present evidence that this is partly due to the increased power of executives over boards – fostering rent extraction by the former, and the greater importance of CEOs on firms performance – increasing the rate of return of CEO's talent.

With respect to job polarisation, technological progress has been a common suspect. For example, Autor, Levy and Murnane (2003) put forward the hypothesis that the widespread introduction of computers has substituted workers in routine-intensive occupations, providing empirical evidence of such polarisation. The emerging literature on automation and robots taking human jobs also shares some of this view, with the anecdotal evidence from the automotive industry, and the available data on industrial robot stocks providing an indirect basis for it (from 66,000 units in 1983 to 1.6 million units in 2015, according to International Federation of Robotics, 2016).

Another important factor behind job polarisation is increasing globalisation and offshoring. There is considerable documentation of the shift of manufacturing production away from developed countries into East Asia, particularly China (Baldwin, 2013). The electronics industry and Foxxconn – Iphone's manufacturer – are a remarkable example of this phenomena. An emerging literature on global supply chains and global production networks is digging into these phenomena, particularly by studying the patterns of trade between countries. As Baldwin (2013) argues, part of the greater interdependence between the production networks of countries has been facilitated by advances in ICT (related to technological change). But it is also due to large wage differences between the "West" and the "East", making the use of global production networks highly profitable. The role of offshoring and import substitution in the fall in wages and employment for workers in advanced economies, particularly in manufacturing, has been highlighted by several studies (e.g. Autor, Dorn and Hanson, 2013; Ebenstein et al., 2014).

0.3 This thesis

This thesis explores the effect of technology, education and skills on the rise and fall of the middle class. In this thesis I use both fixed proportion and variable proportion definitions of the middle class, to provide a more comprehensive understanding of its evolution. Using both theoretical and empirical methods, this thesis advances our knowledge of the role the aforementioned factors play in the development of an economy, the distribution of gains among members of the society, and the labour market adjustment processes that this development entails.

Chapter 1 studies how technological change and education impact the development process of an economy. In particular, it constructs and simulates an overlapping generations model that gives rise to a Kuznets curve. As described above, this is a historically relevant feature of the development process of advanced economies. The dynamic of the model is very simple. Agents in this economy need to decide whether to become skilled or not, by investing in education. Part of this decision depends on the inheritance they receive, which is compared with the cost of education and the financial cost of any borrowing required to pay for it. After this decision, workers join either the unskilled sector, or the skilled sector. The latter sector also uses capital in production. Importantly, there is an externality arising from the aggregate stock of education (or human capital), which only benefits workers in the skilled sector. As such, this model focuses on skill-biased technical change (SBTC). Before dying, agents leave a bequest to their offspring, which depends on their own wealth. As such, there is an intergenerational transmission of bequests within dynasties, which gives dynamics to the model. It is shown that, if the externality is high enough, the economy grows over time, as an increasing number of agents choose to become educated. It is also shown that along this process, inequality increases first, but then falls, as the majority of the population become educated. Finally, a simple analysis of the evolution of the middle class is conducted, in accordance with the methods and definitions outlined earlier. This confirms that along the development process, there is a “rise” in the middle class. This chapter contributes to the literature by providing a micro-foundation of the Kuznets curve, amid a model that is also relevant for studying the conditions that enable development in the long-run.

The aforementioned model is well posed to study long term development trends, like those presented in Figure 0.1. Nonetheless, it does not provide the tools to study job polarisation. It was shown above that this is a pervasive phenomenon in many advanced economies. **Chapter 2** focuses on the erosion of the middle class by means of a model with a more elaborated treatment of the labour market. The focus here is not on long-term development but on changes to employment shares across occupations, coming from both the demand and supply sides of the labour market; namely, from technological change and the evolution of the workforce skill set. To evaluate this, a simple sorting model based on three *tasks* is presented. These tasks are manual, routine and abstract, a categorisation common in the literature. Central in the model is the assumption of heterogeneous ability of workers in carrying out task. In particular, I assume the distribution of ability is log-normal, albeit with different parameters across tasks. This particular distributional choice is supported by the empirical analysis carried out in Chapter 3.

In this model, workers sort into occupations, where they perform a single task. As such, the concept of occupation and tasks become synonymous. The sorting depends on comparative advantage of workers' abilities for these tasks. Each worker contributes to the production of a good according to his or her ability. The total output of workers in a given task is equivalent to the sum of ability of all workers in that task. These outputs are then used by firms as inputs on the production of a final good, which is sold to consumers. In a competitive environment, firms demand inputs from workers, and workers supply inputs to firms. The relative price of these inputs (related to workers wages) adjust in order to bring about the final sorting equilibrium. Competitive markets ensure this output is the maximum possible. After characterising this equilibrium, the model is used to explore the role of both task-biased technical change (TBTC) and changes in ability on job polarisation.

This chapter contributes to the literature by proposing a new channel by which job polarisation can arise, namely changes in ability, a channel which so far has been neglected in the literature. It does so by presenting a simple but useful model where the role of ability can be captured. This model is based on assumptions regarding ability which are empirically relevant, as Chapter 3 shows, in stark contrast with other work, based on assumptions which are at odds with the evidence.[6] I believe that the model is useful in helping understand polarisation in real economies, as Chapter 4 shows.

The particular assumptions about the distributions of ability are central to the model in Chapter 2. **Chapter 3** is entirely dedicated to the empirical identification of these distributions, focusing on the United Kingdom. In particular, the goal is to characterise the complete skill-set of workers, including (i) ability of workers in occupations in which they are observed (e.g. ability of manual workers in carrying out manual tasks), and (ii) ability of workers in occupations in which they are not observed (e.g. ability of manual workers in carrying out routine tasks).

Workers' ability in "observed" occupations (point (i) above) are estimated from a panel-data Mincer equation, using data from the British Household Panel Survey, covering the 1991-2008 period. In effect, a regression approach allows us to control for all other factors known to determine earnings – like education, experience, gender, etc. My hypothesis is that residual variation in wages reflects the underlying ability of workers. This exercise yields an estimate of ability that is positively skewed, qualitatively similar to a log-normal distribution. After this, workers' ability in "unobserved" occupations (point (ii) above) is imputed using the sorting model developed in Chapter 2, together with the results from the Mincer equation. In effect, workers occupational choices act as *revealed preferences*, providing useful information about their ability in occupations in which they are not observed. To provide further support for this imputation method, the central assumptions of the model in Chapter 2 are empirically evaluated. The combination of the two approaches leads to the final result: the distribution of ability for the three tasks is shown to be positively skewed. Qualitatively, a log-normal distribution is the simplest and most well-known among such distributions, justifying the assumptions in Chapter 2.

Chapter 3 provides many contributions to the literature. The most important is that it provides the first attempt to estimate the **complete** ability distributions (i.e. both in "observed" and "unobserved" occupations), using a wage (or Mincer) regression method. Existing literature is incomplete in this respect.

[6] For example, Autor and Dorn (2013) assume that ability of workers in manual and abstract task is homogeneous, and the distribution of ability in routine task is exponential. Furthermore, they assume no mobility of workers between routine and abstract tasks, which Chapter 3 shows to be relevant.

Finally, **Chapter 4** builds upon the model in Chapter 2 and the estimation of ability distributions in Chapter 3 in order to identify the nature of task-biased technical change in the UK economy. The identification method is very simple. First, employment shares for different tasks are observed from the data.[7] Second, the parameters of the distribution of ability are obtained from the results in Chapter 3. Since ability estimation was carried out for several years, the ability distributions can be obtained for different points in time. Any change in these estimates represents a **supply shock** in the labour market. Then, combining observed employment shares and the effects of the supply shock, it is possible to identify the **demand shock**, which in our model corresponds to task-biased technical change. This identification is useful because it means that several counterfactual exercises become now possible. In particular, by analysing how the labour market would have behaved in the absence of one of these changes, we can calculate the relative importance of, on the one hand, changes in ability, and, on the other hand, task-biased technical change in leading to job polarisation, increases in inequality, and the erosion of the middle class. The conclusion is a striking one: that task-biased technical change is the sole driver of these processes.

The key contribution of this final chapter is that it provides an estimate of magnitude of the task-biased technical change affecting an advanced economy, namely United Kingdom, and an estimate of the extent to which this technical change is biased as between manual, routine, and abstract occupations. To the best of my knowledge, this is the first such study. It enable us to quantify the consequences of such technical change for the labour market.

[7] The methods in this and the previous chapter require a mapping between occupations and tasks. This is, there is a need to link every occupation available in occupational code standards (like ISCO or SOC) into one fo the three tasks – manual, routine, or abstract. The method used in this thesis build upon the existing literature on the issue. Section 3.6 in Chapter 3 presents a detailed description of this mapping.

1

Inequality dynamics and economic development: a microfounded Kuznets curve

1.1 Introduction

Which are the conditions that enable economic development? How does inequality and the middle class evolve along this development? This chapter explores these issues through an OLG model of human capital accumulation, where inheritance and an externality in the level of human capital play a central role both in the process and destination of the economy.

Historically, the development of some advanced economies like the US and UK have has been characterised as following a Kuznets Curve (KC). The KC is an inverted U-curve relationship between inequality (y-axis) and economic activity (x-axis), proposed by Kuznets (1955, 1963). The KC has received considerable attention in the literature, and support for it is particularly evident in advanced economies (Milanovic, 2016). Given the historical relevance of the KC, particular attention is presented here in a development process which is consistent with such historical fact. In fact, the model is calibrated and simulated in order to show how such development process arises.

The model combines the overlapping generation (OLG) setting in Galor and Zeira (1993) with an endogenous growth framework, inspired by both Romer (1986) and Lucas (1988).

The OLG framework is relatively standard. The economy is populated with a fixed number of dynasties, each composed of a young and an old agent. Young agents take a vocational choice, by deciding whether to go in education (and become skilled workers when old) or remain unskilled. This decisions depends on the cost of education, the skill premium, the cost of borrowing related to the benefit of lending and, crucially, the inheritance they receive from the old. Heterogeneity in inheritance – the only source of heterogeneity in the model – becomes central in the dynamic of human capital, wages, inequality, and output.

This OLG framework operates amid a small open economy composed of two sectors producing the same good but with different technologies. There is a "traditional" sector, which uses unskilled labour only, and a "modern" sector, which uses both capital and skilled labour. Capital is internationally mobile whereas labour and goods are not. This framework can be rationalised as follows. Initially, an economy starts with no skilled workers; only the "traditional" technology is available. Then, new inventions allow for the introduction of a better way to produce the same good. This "modern" sector requires skilled workers, which are provided through the educational system.

Importantly – and the central departure of this model with respect to that of Galor and Zeira (1993), there is a positive externality affecting the "modern" sector, arising from the number of skilled workers.[1] This is to say, the externality arises from the level of human capital, embedded in skilled workers. This specification follows the spirit of Romer (1986), in that the stock of an "accumulated" factor produces technological externality, which cannot be internalised by firms. Yet, the specification also borrows from Lucas (1988), in that such factor is not physical capital but human capital, and that such accumulation requires to divest resources from production (as the young are outside the labour market). This combination, amid such OLG setting, makes this model unique.

The technological externality can also be thought as skill-biased technical change (SBTC), in the sense that it expands the skill premium, benefiting skilled workers only.[2]

[1] In the baseline model in GZ93, the three factor prices (the return to capital, the unskilled wage and the skilled wage) are exogenous. The introduction of endogenous economic growth endogenises the skilled wage, allowing for a process of human capital accumulation that drives the skill-premium up.

[2] There is no feedback mechanism by which unskilled workers benefit. Future work aims to expand this by introducing a second consumption good which is profitable to produce only with unskilled workers. The elasticity of substitution between such goods is to become a central element in defining the behaviour of the skill premium.

Complementary, this framework represents a process of structural transformation, as the economy transits from a traditional-based production to a modern-based one.

The consequences of this externality are far reaching. In my framework, there are a set of conditions which enables a process in which more and more dynasties to move towards skilled jobs, which expands productivity in the "modern" sector, thereby expanding the skill premium. The build up of wealth among dynasties derives a steady state where the economy fully migrates from the "traditional" technology towards the " modern" one, and every worker is skilled. In other words, the addition of endogenous growth introduces dynamics into the potential stages at which an economy can be.

After describing and solving the model, I use it to simulate a development process of a Kuznets curve, without losing sight that other outcomes are also possible. Starting from a "subsistence" economy based on traditional production methods, a new, modern technology is introduced. Given the externality the latter technology implies, a transition process commences, whereby an unskilled workforce is gradually replaced by a skilled one, as generations evolve. Along this process, wealth accumulation among dynasties follows a pattern where wealth (and wage) inequality increases in the early stages of development – as only the richest among the population benefit from this SBTC. Later on, more and more agents become skilled, leading to higher development and the surge of a "middle class", where inequality starts to decrease. In the end, the economy fully develops, reaching a steady state where bequests equalize and inequality disappears.

This chapter contributes to the literature by highlighting the role of inheritance on development of the economy and the middle class. The benefits of SBTC accrue over more and more dynasties as these are able to accumulate wealth. In effect, two countries facing the same SBTC (e.g. as a new technology becomes available) can benefit differently from it, depending on the nature of the wealth distribution of its members, which is essential in determining the extent and speed of human capital accumulation, necessary to make use of this new technology. Countries where most of the households are poor and just a few very wealthy individuals exist might not benefit from new technologies, as level of human capital needed to "kickstart" the development process and generate sufficient degree of externality to expand productivity (and the skill premium, vital for incentive to become educated) might not be enough. Conversely, in more egalitarian and wealthy societies, adoption of new technologies are more likely to occur.

Similarly, this chapter highlights the role of externalities arising from SBTC, or new technologies. Countries adopting technologies with greater complementarity with human capital will see faster human capital accumulation, and therefore higher growth and ultimately a stronger middle class. Conversely, countries adopting "laggard" technologies (in the sense that the skills workers need to acquire generate little change in aggregate productivity of other skilled workers), will see slower growth of output and middle class formation.

Finally, regarding the Kuznets curve, this chapter provides a new and rather simple micro-founded version of this famous empirical observation, and a more general framework under which to study the complex relationship between inequality and development.

1.1.1 Road map

The Chapter continues as follows. Section 1.2 reviews the empirical and theoretical literature regarding the relationship between inequality and economic development. Section 1.3 presents the complete model of this chapter, starting with the demand (or production) side and continuing with the supply (or agent's decisions) side. Section 1.4 characterises the solution of the model, based entirely on the equilibrium of the skilled labour market, after which all other endogenous variables can be found. Section 1.5 – the most important section – uses the model to study a development process that reproduces a Kuznets curve. Using a narrative, graphical and computational approach, it is shown how the model leads to the desired goal. It also provides a brief description of how the middle class evolves in this economy. Finally, Section 1.6 concludes.

1.2 Literature Review

Income inequality (and inequality in general) is probably one of the most multidimensional topics within economics, reflected by the broad range of disciplines within economics that have contributed to the topic, including development economics, new institutional economics, political economy, business cycle theories, health economics, growth theory, economic history, and so on. This chapter relates in particular to the relationship between income inequality (II henceforth), usually measured as the Gini index, and the level of economic development (ED henceforth), usually measured as the GDP or GDP per capita.

Several questions arise with respect to the relationship between II and ED. Does inequality increase as an economy develop, or does it fall? Actually, is there any relationship at all? What is the direction of causality? What are the factors on which this relationship hinge? Naturally, a description of the literature on all these issues requires a complete survey. Here, I focus on the main lines of thinking, in order to understand where this work fits in.

First, an outline of the papers more related to this chapter is presented. These papers relate to non-monotonic patterns of inequality (like a Kuznets curve) as the economy grows. Then, the focus is on the particular channels by which inequality harms growth. Some of these are used in this chapter whereas others are not. Some of them might lead to nonlinear II-ED relationships, or the link might be linear. In any case, these are important in the literature, reason why these are presented.[3]

1.2.1 Non-monotonic relation between income inequality and economic growth

The Kuznets curve

One of the best known contributions to the understanding of the link between II and ED were the seminal works by Kuznets (1955, 1963), in which by a mix of empirical

[3]Certainly, many works consider inequality irrelevant for understanding economic growth. Most of the neoclassical growth literature, either with exogenous or endogenous growth, is based on the representative agent paradigm, where inequality is irrelevant. Also, notice that a model with heterogeneous agents (or heterogeneous firms) is not sufficient to turn inequality relevant for growth. Under certain assumptions, a model with heterogenous agents can be represented as one with a representative agent (Caselli and Ventura, 2000; Krusell, Smith and Jr., 1998).

data and theoretical analysis, he advanced that inequality had followed an inverted U-curve (usually called the Kuznets curve) for advanced countries. In his account, this occurs as the economy transits from an agricultural and rural economy (where inequality is lower) to an industrial and urban economy (where inequality is higher), and then as the economy keeps growing, inequality decreases because the non-agricultural low income groups gain political power which results in institutional change that favours them (thus, inequality was a by-product of industrialisation). This dualist vision of the development process was previously analysed in a seminal paper by Lewis (1954), whose model can also be used to derive a Kuznets curve. Regardless of the limitations of Kuznets' original argument and evidence (that he was aware of), the hypothesis anchored deep into the profession, receiving huge empirical attention since then (see Moran, 2005 for a historical account of the impact of the Kuznets curve hypothesis in the social sciences in general).

In terms of the empirical evidence the Kuznets curve has received mixed support. Just to mention a few recent studies, Barro (2000, 2008) finds that the Kuznets curve does exist for a group of countries (between 54 to 92 countries depending on the period used) even when inequality is measured in three different ways. Barro also finds that the curve has been relatively stable, although more pronounced in recent years. Banerjee and Duflo (2003) also support the Kuznets curve story while stressing the errors generated when imposing a linear relationship between II and ED (as in Forbes, 2000 and Li and Zou, 1998, both finding no evidence of the Kuznets curve). Higgins and Williamson (1999) finds evidence of a Kuznets curve for 85 countries for a period of 40 years while also stressing the important role that the cohort size has on inequality, a factor usually not considered in previous studies. Others argue that the Kuznets curve is reflecting the "Latin America effect" (the region with usually the highest levels of inequality). Deininger and Squire (1998) show that controlling for that effect eliminates the Kuznets curve. Gallup (2012) finds evidence of a Kuznets curve from a cross-section perspective, but shows that the relationship is the converse – a U-curve – when a time series approach is taken. His methodology is interesting because instead of focusing on years, the time-series perspective uses GDP levels as a measure of development stage, as the original theories suggest. More recently, Milanovic (2016) compiles data for the US and

UK, showing a clear Kuznets curve between at least the 17th century, with peak around 1860, and bottom in the post-war period.[4]

These contradictory results might be partly due to methodological differences or poor data quality but also are related to how complex the relationship between II and ED is. While for some countries it could be best understood as a Kuznets curve, for others it might not. As the same Kuznets argued the empirical observation was space-time dependent.

Micro-foundations of the Kuznets curve

While the empirical literature testing the Kuznets curve is large, its theoretical micro-foundation (like the one develop in this chapter) has been less studied. As already mentioned, Lewis (1954) model of development as a process of migration of workers from a subsistence sector (usually associated with agricultural and rural sectors) to a capitalist sector (industrial or urban sectors) results in a Kuznets curve since at the beginning the migration of workers imply that capitalists increase their profits with wages remaining constant whereas a the end, once all workers have migrated, wages start to increase, reducing inequality. Kuznets (1955, 1963) argument was based on this dual development process.

Rauch (1993) formalizes Lewis and Kuznets hypotheses by means of an OLG model of progressive urbanization with two sectors (agricultural and urban) but differencing between an urban formal subsector and an urban informal subsector. Thus, at early stages agricultural wage is low due to high rural population, resulting also in low inequality since most of the population is in that sector. As population starts to migrate to the urban sector inequality increases because this sector is more unequal (the formal subsector pays a higher wage than the informal subsector). Along this process the wage in the agricultural sector increases. Finally, as most of the population is in the urban sector the rural wage is as high as the urban formal sector wage, hence reducing inequality.

[4] The author goes further to argue that there is not just one Kuznets curve, but Kuznets waves, with the recent increases in inequality reflecting a new, current wave.

There are other models that depart from this dual perspective and yet can also generate a Kuznets curve. For instance, Galor and Tsiddon (1997a) propose an OLG model without bequests nor financial frictions where agents have different abilities and parent' human capital level, both affecting their own human capital level, income, and consumption (thus, there is an externality in education, like in my model). Production of a single good is done in several industries, each with endogenous growth modelled as "learning-by-doing". Thus technological progress and economic growth depend on the initial distribution of income and abilities. The model results in that in early stages human capital concentrates in a few advanced sectors, concentrating income and thus rising inequality but the same higher growth increases return for labour in other sectors, thus increasing human capital in general, reducing differences in income and inequality. In a related paper, Galor and Tsiddon (1997b) obtain a similar result but assuming that agents have equal abilities to become educated. The interplay between the "local" externality in the education of agents' parent (which directly affects these agents) and the "global" externality in the education of agents' parent (which accelerates growth and wages) generates the inverted U-curve in inequality.

Glomm and Ravikumar (1998) show how the Kuznets curve arises under the presence of increasing returns in the accumulation of human capital by agents. These increasing returns are beneficial at the beginning but because agents are constrained by their time endowments human capital cannot grow further (thus, generating what they call short-term increasing returns to scale). Although their model has no bequests, inequality dynamics come from an externality in human capital, where agents education is a function of parents' education (thus parents' education is a proxy for bequests). The model contrasts with the one in this chapter because here increasing returns are always operating (in their terminology, there are long-run increasing returns to scale).

Grossmann (2008) builds an OLG model with divisible human capital and bequests, close to the model in Galor and Moav (2004) but without financial frictions. The key mechanism generating the Kuznets curve is the interplay between the increasing marginal propensity to save (motivating rich agents to leave more bequest than poor agents, so increasing inequality) and the decreasing marginal return of human capital (motivating poorer agents with less human capital to invest in education and so reducing inequality).

Other papers stress the role of an imperfect financial market in generating a threshold that separates those that participate in this market from those that cannot. For example, Greenwood and Jovanovic (1990) present a model with information imperfections in the allocation of resources into production by agents. Financial intermediaries help agents to overcome these imperfections and thus those agents that invest through these intermediaries obtain higher returns than those that do not use them. Yet, there is a fixed cost of using financial intermediaries so only those rich enough can use them. The key of the model is that these financial intermediaries vary endogenously. Initially the financial market is poorly developed, only some agents invest, the economy grows slowly, and inequality is low. As the financial market develops along economic growth, more agents participate in the market, and inequality increases due to the polarization of agents. At some level most agents use the financial market and inequality starts to decrease until in the steady state all agents behave equally.

Aghion and Bolton (1997) focus on the role of the financial market when the interest is determined endogenously (not exogenously like in here). In a model without human capital, agents need resources to invest in profitable projects, increasing their wealth, which then "trickles-down" by increasing the amount of resources available the next period, hence allowing poorer agents to invest in these profitable projects. In the long run all agents end up investing and being rich, but with a middle phase of higher inequality.

From an institutions perspective, Acemoglu and Robinson (2002) also obtain the Kuznets curve by modelling how inequality increases with the growth of a capitalism mode of production that benefits the elites but then decreases as poor masses demand for redistribution policies. Perotti (1993) also focuses on the political equilibrium in terms of the outcome of taxes to income. Taxes are redistributed enabling poorer agents to invest in human capital. Just like this chapter, there is an externality of acquisition of skills that spills over to the rest of agents, thus leading to a long run when all agents are educated and rich, but with a middle phase of high inequality due to only a fraction of the population being rich. Yet, in their model there is no financial market.

There are other papers that are able to generate a non-monotonic relationship between II and ED but not exclusively the Kuznets curve. García-Peñalosa and Turnovsky (2011) develop a Ramsey model with an elastic labour supply and two types of heterogeneity (initial capital ownership and time-invariant labour skills) in which inequality is affected by the leisure-work decisions and the change in the relative value of factors, leading to a complex link between II and ED that depends on the particular assumptions. Borissov and Lambrecht (2012) combines both approaches of indivisible human capital levels (skilled versus unskilled) with divisible levels within the skilled sector in a model with endogenous growth where unskilled workers benefit from the aggregate level of capital, showing that the II-ED relationship can be increasing, decreasing or that of a Kuznets curve. In these two papers there is no room for financial frictions.

1.2.2 Channels through which inequality and economic development interact

The previous section presented models leading to a Kuznets curve. Another perspective on the issue relates to the channels by which inequality affects growth. Some have been mentioned already, others can be added – which do not necessarily lead to Kuznets type of development. Here a brief review is presented. For a more comprehensive analysis, see Ehrhart (2009) and Galor (2009).

One of the causality channels most studied in the literature relate to credit market imperfections. This is a credit market that restrict access to funds for some agents (Banerjee and Newman, 1993, Piketty, 1997) or that implies higher costs for poorer agents (GZ93). Any of these might have detrimental effects for growth because not all agents develop their full potential as skilled workers (human capital underinvestment), thereby delaying or halting the transition from low to high productivity technology; the economy ends up in an suboptimal equilibrium. This hypothesis has been generally supported by empirical studies either studying the direct effects of credit constraints on the economy or their indirect effects through human capital (Deininger and Squire, 1998, Khalifa and Hag, 2010, Papageorgiou and Razak, 2009, Perotti, 1996). Notice that if the credit market is perfect (in the sense that the borrowing and lending rates are the same), then inequality would have no effect on growth. On the contrary, when there is an imperfection, inequality might slow down or even halt growth completely. This channel is part of the framework used here.

Another factor that links inequality and growth relates to endogenous fertility decisions. While the general modelling setting is to assume a fixed population (like here), some papers explore the consequences of endogenous fertility. For example, agents might face a trade-off between "quantity" and "quality" of children education (DeNavas-Walt, Proctor and Smith, 2012, Kremer and Chen, 2002, Morand, 1999, Perotti, 1996). For example, unskilled workers are more likely to have more children because their opportunity cost (i.e. wage) is smaller. This might increases the size of the unskilled generation in the next period. If the unskilled wage is endogenous, this might lower the unskilled wage even further, fostering a poverty cycle. This channel is not explored here, but could be certainly incorporated, such that bequests are divided over the number of children. This would certainly lower the extent of social mobility and might even stop or reverse the development process.

Political power is also a channel of relevance, in particular with relation to the extent of redistribution. For example, the distribution of wealth seem to have a direct impact on the type of policies chosen, some more pro-growth than others (Alesina and Rodrik, 1994, Bourguignon and Verdier, 2000, Easterly, 2007, Galor, Moav and Vollrath, 2009, Gradstein, 2007, Persson and Tabellini, 1994). This could be introduce in our model through endogeneity of the educational cost, or access to loans, or taxation. Still, at the moment this is not accounted for. Another explored channel relates to social instability. Some authors see the negative effects of inequality coming from the social and political volatility that inequality might generate. (Alesina and Perotti, 1996, Dutt and Mitra, 2008, Roe and Siegel, 2011). This instability can be reflected in higher interest rates – perhaps a premium with respect to the international, risk free rate. This has negative consequences on optimal capital levels.

In conclusion, there are many frameworks under which a Kuznets curve might arise. Some of these use financial market imperfections and/or human capital externalities generating endogenous growth, features that have also received empirical support. This chapter uses both in a rather simple but new framework, that can derive a Kuznets curve too (among other possible dynamics). Thus, this model contributes to the literature by providing a new micro-foundation of the Kuznets curve hypothesis. Due to time constraints, the model cannot be extended with more realistic features but this is expected to be part of my future research.

1.3 Model[5]

The economy has one good, which can be produced with two technologies. The "modern" one uses both capital and skilled labour, whereas the "traditional" one uses only unskilled labour. As such, this economy is composed of four markets, namely the final good, skilled and unskilled labour, and the financial market for capital/loans. The latter is trivial in the sense that there is open access to capital markets, and the international rate is fixed (the economy is small). Thus, any national shortage (excess) of demand with respect to locally raised funds is compensated with foreign capital inflows (outflows). The equilibrium in this market is then independent on the other three markets. The two labour markets are of central interest in the model. Embedded in an OLG framework, young agents decide to remain unskilled or to become skilled by investing in education. This decision depends on their inheritance, and the skill premium – i.e. the return to human capital investment. As such, both labour markets are cleared simultaneously, since there is just one choice young agents have to make, namely to be unskilled or skilled. Once agents make their decisions, these two markets are cleared. Finally, by Walras' law, the goods market is also cleared. This market is also trivial because there are no trade-offs involved, as the good is unique. The price of this good is the numéraire to which wages are relative.

In consequence, to characterise the model's complete solution it is enough to focus on the equilibrium in one labour market only. Since there is an externality in the "modern" sector arising from the stock of human capital in the economy, all the action of the model is happening in the skilled labour market. Therefore, the presentation of the model puts particular emphasis on that market.

[5]For greater comparability with Galor and Zeira (1993), the notation here follows theirs closely. Also, throughout the rest of this chapter those variables marked with an asterisk denote an optimal level while variables without a time subscript are either constant or at their steady state, depending on the context.

First, the production sector is studied, from where the demand for skilled labour is obtained. Second, agents' decisions are studied, from where the supply of skilled labour derives. The equilibrium is studied in the next section.

Importantly, the model is analysed from the perspective of agents and firms. Because of the externality in one sector, their private decisions are suboptimal with respect to a central planner perspective.

1.3.1 Production

General setting

The single good of this economy can be produced in two technologies or sectors. The "traditional" sector (denoted with a superscript n) uses solely unskilled labour, whereas the "modern" sector (denoted with a superscript s) uses physical capital (capital henceforth) and skilled labour. Aggregate production (or GDP) is equal to the sum of the production in the skilled and unskilled sectors. The price of this homogeneous good is the numéraire in this economy, meaning all monetary variables in the model are in real terms.

Each sector is populated by a certain number of homogeneous firms operating with constant returns to scale (CRS) technologies in a competitive market (price takers, zero profits, no frictions). Under these conditions there exists an aggregate production function for each sector equal to the sum of the production function of individual firms.[6] Thus, there is no need to focus on individual firms – which number is irrelevant for the solution. The aggregate solution maps directly onto individual firms' ones.

Finally, the economy is small and open to the international capital market, which sets the internal interest rate at the world rate, r. There is no trade in goods nor migration flows. There is no government.

[6]See Felipe and Fisher (2003) for a complete discussion on the conditions under which an aggregate production functions can be built from the individual production function of firms.

Traditional sector

The aggregate production function in this sector is:

$$Y_t^n = w^n L_t^n \tag{1.1}$$

where L_t^n is the unskilled labour at time t, and w^n is the aggregate productivity level in this sector, which is positive, constant, and exogenous.

The aggregate demand for unskilled labour is derived from the marginal product of L_t^n. In this competitive market, it equals the unskilled wage:

Demand for unskilled labour

$$w_t^n = \frac{\partial Y_t^n}{\partial L_t^n} = w^n \tag{1.2}$$

Notice this demand is perfectly elastic and constant.

Modern sector

Each firm has a Cobb-Douglas production function with Constant Returns to Scale (CRS). The aggregate production function in this sector is:

$$Y_t^s = A_t^s {K_t}^{\theta} \left(L_t^s\right)^{1-\theta} \tag{1.3}$$

where K_t and L_t^s stand for capital and skilled labour at time t, respectively. As usual, capital is defined as a stock at the beginning of the period, and accumulates based on agents savings.[7] θ is the elasticity of capital, bounded within the unit interval $(0, 1)$. The technological factor A_t^s, common to all firms, accounts for the aggregate level of knowledge available at period t. It is defined as:

$$A_t^s \equiv A^s \left(L_t^s\right)^{\psi} \tag{1.4}$$

[7] Since the economy is open to international capital markets, the dynamic of capital has no effect on the model. Capital could fully depreciate and the model would still reach an equilibrium, but with a different balance of payment equilibrium. See Appendix A for more details.

where A^s and ψ are positive, time-invariant parameters. The first is a scaling factor while the second represents the size of the externality. If $\psi = 0$, this model reduces to that in GZ93.

This specification implies that endogenous growth comes from an externality generated by the level of skills (or human capital), which spills over to the rest of the economy by increasing the aggregate pool of knowledge A_t^s. Higher knowledge implies higher productivity for all factors (and firms) within the sector. Yet, this aggregate level of knowledge is a non-rival public good and so the benefit is not internal to the firms but external to them. The result is increasing returns to scale for the industry but constant returns to scale for each firm, which implies that perfect competition holds in this sector. As a result the private outcome coming from firms' optimisation does not lead to the social optimum.[8]

The specification in equation (1.4) is related to the "accumulation of knowledge" approach in endogenous growth theory. The origin of this perspective comes from Arrow (1962), who argued that firms' investment leads to a process of "learning-by-doing", which increased the level of knowledge in the economy, increasing the productivity of workers and preventing the marginal product of capital to fall. Romer (1986) is an important example of this approach to endogenous growth theory. Yet, unlike here, Romer's approach is based on the accumulation of knowledge from physical capital. As such, the source of endogenous growth is capital accumulation. In our case, the source of growth is the stock of human capital, making this model closer in spirit to Lucas (1988), which also includes the level of human capital in the economy as a factor affecting workers and capital productivity. Additionally, as in Lucas (1988), here the "production of knowledge" requires to divert resources from production, as young agents cannot be both at work and in education. This externality can also be thought as a process of SBTC, which affects skilled workers only. Similarly, it can be seen as a process of structural transformation, where a new, modern technology becomes available, perhaps gradually replacing traditional production methods.

[8] As shown later, the private demand for skilled labour is lower than the social demand for skilled labour so firms demand less skilled workers than socially optimum, reducing welfare whenever not all agents are skilled.

Combining equations (1.3) and (1.4) yields:

$$Y_t^s = A^s K_t{}^{\theta} \left(L_t^s\right)^{1-\theta+\psi} \tag{1.5}$$

This reflects a standard result in the endogenous growth literature focusing on externalities, namely that the sector as a whole presents increasing returns to scale (IRS) even though each firm operates under CRS.

Analogous to the case of unskilled labour, using the aggregate production function to derive the aggregate demand for capital is a valid procedure because capital does not present an externality (although not shown, the proof is simple). The market is competitive in this sector so the marginal product of capital is equal to the net rate of return on capital. Moreover, capital is the only asset of the economy so the net return to capital is equal to the world interest rate. Therefore:

$$r = \frac{\partial Y_t^s}{\partial K_t} = A^s \left(L_t^s\right)^{\psi} \theta \left(\frac{K_t}{L_t^s}\right)^{\theta-1} \tag{1.6}$$

$$= A^s \left(L_t^s\right)^{\psi} \theta k_t{}^{\theta-1} \tag{1.7}$$

From here we can derive the aggregate demand for capital *per skilled worker*:

Demand for capital *per skilled worker*

$$k_t^{D^*} = \left(\frac{A^s \left(L_t^s\right)^{\psi} \theta}{r}\right)^{\frac{1}{1-\theta}} \tag{1.8}$$

The fact that the capital demanded per skilled worker is increasing in the level of skilled labour is a by-product of the externality. The higher the stock of knowledge in the economy, the more productive this capital is. This is a feedback effect that creates a virtuous cycle because, as seen below, the demand for skilled labour is also increasing with capital. This cycle allows the economy to grow over time, contrasting sharply with the specification in GZ93, where the optimal level of capital is fixed (e.g. setting $\psi = 0$ in equation 1.8).

The skilled labour market is competitive in the sense that firms make no profits and take prices as given. Moreover, firms do not internalise the externality. Thus, when obtaining the aggregate demand for skilled labour, it is important to hold A_t^s as fixed. Thus, this demand is obtained from $w_t^s \equiv \left(\frac{\partial Y_t^s}{\partial L_t^s}\right)_{A_t^s}$, which is equivalent to aggregate each firm's demand over w_t^s. The result is:[9]

Partial demand for skilled labour

$$w_t^s = A^s \left(L_t^s\right)^{\psi-\theta} (1-\theta)\, {K_t}^{\theta} \tag{1.9}$$

Importantly, equation (1.9) is a "partial equilibrium" demand in the sense that **it holds for a given level of capital in the economy**. Yet, capital demand is in fact a function of the skilled labour, as equation (1.8) shows. Plugging the latter into equation (1.9) yields the "total equilibrium" aggregate demand for skilled labour:

Total demand for skilled labour

$$w_t^s = A^s \left(L_t^s\right)^{\frac{\psi}{1-\theta}} (1-\theta) \left(\frac{A^s\theta}{r}\right)^{\frac{\theta}{1-\theta}} \tag{1.10}$$

This demand already incorporates the optimal level of capital demanded by firms. Whenever skilled labour changes, optimal capital also changes, which then feeds back into the skilled labour demand. This feedback is incorporated into equation (1.10), meaning it is not a *ceteris paribus* demand. Yet, for determining the final equilibrium in the skilled labour market, this is the one that matters.

Moreover, whereas the slope of the "partial" aggregate demand depends on the relative values of ψ and θ, **the "total" demand is always upward slopping** (see Appendix B).[10] Albeit odd, this is an outcome of the virtuous cycle created by the positive

[9]Recall this is the **private** demand for skilled labour. The social aggregate demand optimises over A_t^s too. For reference, this is $\frac{\partial Y_t^s}{\partial L_t^s} = A^s \left(L_t^s\right)^{\psi-\theta} (1-\theta+\psi)\, {k_t}^{\theta}$. As can be seen, the difference between the private and the social aggregate demands is the term ψ inside the parenthesis in the latter. Notice that a positive externality ($\psi > 0$) means the privately optimal skilled wage is lower than the socially optimal one, *ceteris paribus*. This is, firms demand a lower level of skilled workers than socially optimal, leading to lower output too.

[10]The curvature of the "partial" demand for skilled labour – equation (1.9) – depends on the value of the exponent of the labour variable ($\psi - \theta$). The latter expression can be positive, zero, or negative, resulting in an increasing, horizontal, or decreasing demand respectively. In effect, any change in the skilled labour generates two contradictory effects in the productivity of skilled labour: (i) a fall in labour productivity due to the decreasing marginal returns to skilled labour, measured by the coefficient θ; (ii) an increase in labours productivity due to the increase in the stock of knowledg ein the economy – the positive externality, measured by the coefficient ψ. The relative value of the coefficients determines the net effect. If the externality is large enough so that $\psi > \theta$, an increase in the quantity of skilled labour demanded will raise the skilled wage; i.e. an upward slopping demand. If the externality totally offsets the decreasing returns ($\psi = \theta$), the demand for skilled labour is perfectly elastic. If the externality is relatively low ($\psi < \theta$), the aggregate demand for skilled labour is downward slopping.

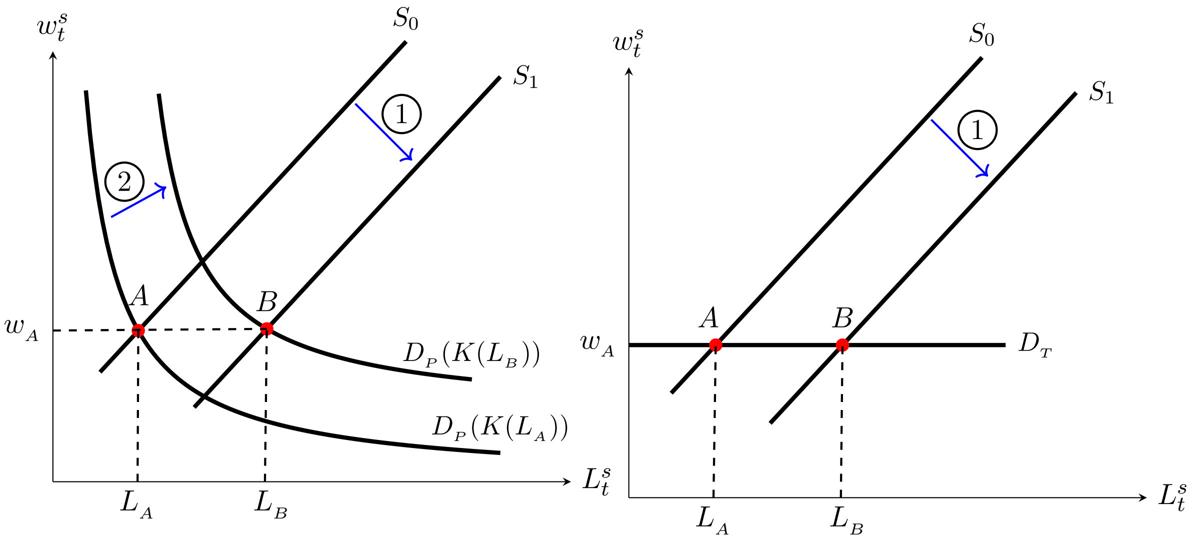

(a) Partial aggregate demand (D_P), for different capital demand levels ($K(L^s)$)

(b) Total aggregate demand (D_T)

Figure 1.1: Response of the partial and total equilibrium demand to an increase in skilled labour supply, from S_0 to S_1, for $\psi = 0$

externality. This upward slopping demand will allow the skilled wage to grow over time, a key property of the solution, to be analysed in details later.

It is insightful to show the differences between the "partial" and "total" demands by presenting the market for skilled labour. The reader might be aware that the supply of labour has not been introduced yet. For the sake of this exercise, let us assume an arbitrary supply. First, let us explore the case of an economy without externality. This is, $\psi = 0$, and the production function reverts to a standard Cobb-Douglas. Panels (a) in Figure 1.1 shows the downward slopping "partial" aggregate demand for skilled labour, together with an arbitrary supply. The economy is initially at point A, with skilled labour equal to L_A and capital stock of $K(L_A$. Now, in a given period, there is an exogenous increase in the supply of skilled labour, from S_0 to S_1. In this economy without externality, the optimal capital *per worker* in the economy is fixed at $k*$ (see equation 1.8). As such, the optimal response by firms is to increase capital from $K(L_A$ to $K(L_B$, level at which capital per worker is again equal to $k*$. In terms of the graph, the "partial" demand would rise, crossing the supply at the same wage than initially (w^A), at point B. This constancy in the wage can be easily seen in equation (1.10), by setting $\psi = 0$. Recall that in a Cobb-Douglas production function, the marginal product of any factor depends on the relative intensity of factors, which in this case is constant.

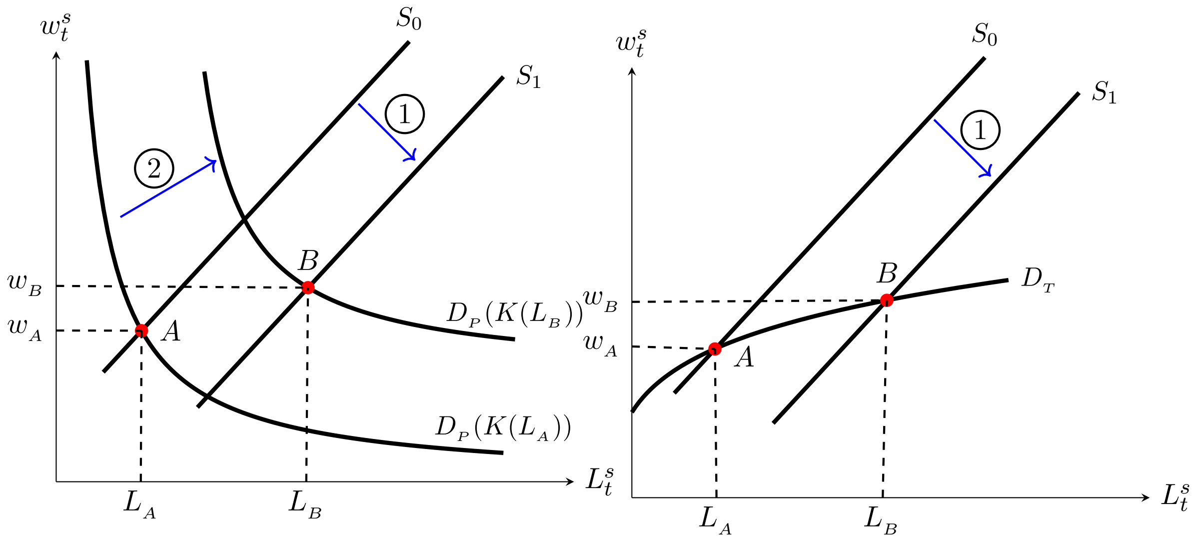

(a) Partial aggregate demand (D_P), for different capital demand levels ($K(L^s)$)

(b) Total aggregate demand (D_T)

Figure 1.2: Response of the partial and total equilibrium demand to an increase in skilled labour supply, from S_0 to S_1, for $\psi > 0$

Now, according to our definition, the “total” demand for skilled labour already incorporates the change in capital. This demand is represented in Panel (b) of Figure 1.1. As shown, it is perfectly elastic, equal to w^A. This demand is constructed by joining the points A and B. More precisely, it is constructed as in equation (1.10).[11]

Let us now move to the case with an externality. This is presented in Figure 1.2. Again, the economy starts at point A. Then, there is an exogenous change in the supply of skilled labour. Just like before, there is an increase in the capital stock. Unlike before, this increase raises the optimal capital *per worker* in the economy. This is because the externality arising from the higher knowledge stock (or human capital stock) feedbacks into greater productivity of capital, fostering a greater investment than in the case of $\psi = 0$. Panel (a) shows how this higher capital per worker increases the skilled wage. Finally, the “total” aggregate demand can be seen in Panel (b). It should be now clear how the externality leads to an aggregate demand that is always upward slopping.[12]

[11]As in GZ93, the skilled wage is in fact exogenous in the model.

[12]The analysis so far is based on a rather loose treatment of inter-period dynamics. Since the full model has not yet been introduced, a more comprehensive analysis is not possible. The process by which the equilibrium is achieved is fully explained in Section 1.4. The reader is invited to revisit this analysis after that section, if this exercise has not been convincing enough.

Some justification for the assumptions in the production side of the model

Economic growth in this model is based on a human capital externality. Whilst most of the evidence support the existence of some degree of human capital externalities, the evidence is mixed when it comes to the actual size of this externality.[13] For example, Acemoglu and Angrist (2000) estimate *local* externalities by computing the effect of average education at the state level on individual wages among a sample of white US citizens, between 1960 and 1980. They find an effect between 1 and 3%, which they classify as modest. However, as they state, their calculations cannot capture *aggregate* externalities, e.g. the effect of ideas generated by high-skill workers that enhance productivity in other parts of the country. This is surely a very important type of externality, and implicit in my model. Winters (2013) studies US metropolitan areas, finding that areas with higher human capital increase the likelihood of local resident in finding a job. Notably, albeit the benefit accrues to workers of all educational levels, the benefit is higher for less skilled workers (against the assumptions of this model). Using data on Italy, Dalmazzo and Blasio (2005) shows that human capital not only produces "production externalities" (i.e. increases productivity of other workers) but also "consumption externalities" (e.g. higher variety of consumer goods and higher quality of public services). There is a large literature exploring other non-pecuniary effects of human capital like crime (e.g. Lochner and Moretti, 2004), or democracy and political stability (e.g. Milligan, Moretti and Oreopoulos, 2004), factors which in the medium term can benefit productivity across the country.

Unfortunately, measuring aggregate externalities is very difficult, because of all sort of endogeneity and identification issues (Moretti, 2004). One attempt, which (tries to) takes care of these issues is Arteaga Cabrales (2010). They use FE and IV methodologies to estimate human capital externalities for 60 countries in the period 1980-2000. Their approach is to estimate the degree of returns to scale in a production function with high-skill and low-skill labour (and capital), finding evidence of increasing returns to scale – between 1.7 and 2.2.

[13] For an example of a paper which finds no external effects of human capital, see Ciccone and Peri (2006).

There are caveats in the above discussion. The model developed here is pertinent for an economy transiting from agriculture to manufacturing. As such, this is intended to represent the historical experience of advanced economies and not their most recent experiences. Therefore, estimates of recent externalities might not provide an indication of past externalities. Unfortunately, to the best of my knowledge, there is no study evaluating the role of externalities in the structural transformation of modern economies.

Another essential assumption of the model is the fact that there is no channel by which this externality affects the unskilled wage, which was assumed fixed. This is, the productivity of workers in the traditional sector does not depend on knowledge. On the one hand, this is as mere simplification that facilitates the solution of the model. As it will be clear later, the choice made by agents regarding whether to become skilled or not depends on the expected skilled wage, which itself is a function of the expected number of individuals investing in education, which again is a function of the expected skilled wage. This is essentially an exercise of matching future demand and supply for skilled workers. Endogenising the unskilled wage via a productivity spillover across sectors add further complication to this setting, as the equilibrium of both the skilled and unskilled labour markets would have to be jointly determined. As Section 1.4 later shows, under the current specification, the equilibrium in the unskilled labour market is simply a residual; solving the skilled labour market is enough to pin down all endogenous variables.

An alternative approach to the productivity spillover mentioned above, which also endogenises the unskilled wage, is to differentiate the goods produced with the two technologies. This is, instead of assuming a unique good, each technology could refer to a different goods, with a positive degree of utility substitution in consumption between them. This way, individuals can still demand the good produced with low skill workers, which through changes in relative prices might induce real wages in the unskilled sector to rise, halting the complete specialisation process (if there is no production, the price would go to infinity). However, this channel also complicated the solution of the model. As Section 1.4 shows, the output market, just as the unskilled labour market, is "residual". This is, as all other markets are in equilibria, the output market necessarily is too (due to Walras Law). Adding a second consumption good would mean having four, deeply intertwined markets (labour and good).

The key question is then how much insight would such extensions add to this framework. First of all, under non-trivial assumptions, they would yield incomplete specialisation. This is different from the case studied in this chapter, where everyone becomes skilled. Some degree of inequality would persists in the steady state too. Yet, as in here, there is no dynamics in the steady state. In the end, a model that can yield a Kuznets curve with or without complete specialisation might not be qualitatively very different.

As such, the model of complete specialisation here is related to the literature on early structural transformation, whereby an economy transits from a mainly agriculture-based society to a manufacturing-based one, as Lewis (1954) and Kuznets (1955), among others, emphasised. For example, Herrendorf, Rogerson and Valentinyi (2014) show how agricultural employment in advanced economies fell from around 80% in 1800 to roughly 5% in 2000, with manufacturing and services – both knowledge intensive sectors – accounting for the increase. Albeit there is no total obsolescence of unskilled labour-dominated production (many of which are nowadays in the service sector), this model's characterisation is a rough match to the aforementioned transformation trend. Being aware of the limitations the current formulation presents, I expect in the future to move towards a more rich set of assumptions.

1.3.2 Dynasties and agents

Having characterised the production market, and the demand for skilled labour in particular, now the focus is on deriving the supply of skilled labour. This comes from the OLG structure of the model, which is very standard.

Dynasties

The economy is populated by N dynasties, indexed by $j = 1 \dots N$. Every period, dynasties are composed of one young agent (y) and one old agent (o). There is no population growth, namely each agent has only one child. The total number of agents in the economy is fixed at $2N$.

Given this notation, the indexation triplets $Z_t^{y,j}$ and $Z_t^{o,j}$ fully identify a particular variable Z with a unique agent within a certain dynasty for a certain period. Yet, to increase exposition's clarity, the j index is generally omitted. The index y and o already distinguishes individual level variables from aggregate ones.

Agents' preferences

Agents derive utility from consumption when old, and by leaving a bequest to their offspring. For simplicity there is no consumption in the first period, meaning that young agent's consumption is incorporated into old agent's consumption. Utility is derived from the *level* of bequests, a specification called imperfect altruism or "joy of giving", a common approach in both OLG and life-cycle models.

The lifetime utility function is a linear combination of both consumption (c_{t+1}) and the bequest left when old ($x_{t+1}^{o,t}$), each expressed in logarithms:

$$U\left(c_{t+1}^o, x_{t+1}^o\right) \equiv \alpha \ln(c_{t+1}^o) + (1-\alpha)\ln(x_{t+1}^o) \tag{1.11}$$

where $\alpha \in (0,1)$, governing the agent's relative preference for consumption over bequests.

Agents are homogeneous in terms of preferences (i.e. α is not indexed to j). The only source of heterogeneity in the model comes from the bequest they inherit, which might vary over time. In other words, heterogeneity is not structural to the model.

Bequests

In every period each old agent leaves a bequest to the young agent within her dynasty. Denote the bequest *received* by a young agent in period t as x_t^y, and the bequest *left* by an old agent in period t as x_t^o. Because there is no tax to inheritances, inflation, or any loss in the bequest process, $x_t^y = x_t^o$.

From the aggregate perspective, the distribution of bequests received by the young generation in period t is denoted as B_t, with probability mass function $b_t(x_t^y)$, and domain $\Re_0^+$.[14] For example, the number of agents with a bequest between 0 and c in period t is given by:

$$\int_0^c b_t(x_t^y)\, \mathrm{d}x_t^y \tag{1.12}$$

The existence of bequests is essential for the model because (i) they add heterogeneity to the model, and (ii) they generate persistence of inequality.[15]

Educational choice

Agents in this economy have two options: (i) do not invest in education when young, remaining unskilled through their life, and therefore working in the "traditional" sector when young and old; (ii) invest in education when young, leaving them outside the labour market during that period, but allowing them to work in the "modern" sector as skilled workers when old. Education costs h, which is constant and assumed wasted in the process. The justification behind a constant education cost is given once agents' optimisation is introduced. Education is indivisible in the sense that agents are either skilled or unskilled. All agents have the same ability for becoming educated. The only constraint is financial, as explained later.

Before moving further, a cautionary note on terminology should be mentioned. In order to simplify the exposition, throughout the chapter the concepts of "investment", "to invest", and so on refer to investment in education unless otherwise stated. Since education is the key element in the model (and not firm's capital investment) this terminology should not result misleading.

[14] For simplicity I assume that all bequests are positive. Yet, the effect of allowing for negative bequests (a debt) in the initial period only delays the convergence toward the steady state.

[15] It is trivial to show that without bequests, any heterogeneity among the old generation will be lost because dynasties' heterogeneity will not be reproduced into the young generation. In this case young agents will make the same decision – either borrow to finance their education or work as unskilled. Bequests turns an otherwise homogeneous generation of young agents into a heterogeneous group of inheritors.

1.3.3 Domestic financial sector

The financial sector in this economy is an intermediary between lenders and borrowers. The supply of funds comes from agents with positive savings, whereas the demand for funds comes from relatively poor agents wishing to invest in education, and from firms in the modern sector wishing to use capital for production.[16]

Financial intermediaries can enforce the lending contracts made with firms without costs. In consequence, firms will never default, making their borrowing rate equal to the prevailing market rate, r. In contrast, there is a positive cost of tracking borrowing agents. In consequence, they have an incentive to default, which creates a wedge between the borrowing rate (i) and the savings rate (r) for agents, wedge that covers this tracking cost.

Figure 1.3 helps to understand how the financial sector works. Blue arrows represent first period flows while red arrows represent next period return flows. Agents lend $a + d$ to the intermediaries. Of these, a go to firms (equivalent to the capital stock of the next period, K^s_{t+1}), and d go to other agents. Firms will borrow abroad in the case that they need more resources (K^F_{t+1}).[17] Next period payments are resolved. Agents that borrowed pay $d(1 + i)$, firms pay $a(1 + r)$, agents receive $(a + d)(1 + r)$ for their savings, and the remnant are wasted as costs of monitoring to ensure that debtors do not default. No resources are left.

1.3.4 Agents' optimisation

Agents make their lifetime decisions when young, by maximising their utility given by equation (1.11). Importantly, the budget constraint of the young agents depends on both her inheritance and the educational choice she makes. In particular, it depends on whether the agent invests in human capital or not (tags **I** or **N**), and if she invests, on whether she is a lender or a borrower (tags **L** or **B**). Since an agent that does not invest and borrows is not possible (recall there is no consumption when young), there are three

[16] As common in these type of models, the final good in this economy has multiple uses: it can be consumed, saved, and used to produce as capital.

[17] See Appendix A for a complete analysis of capital and the external accounts.

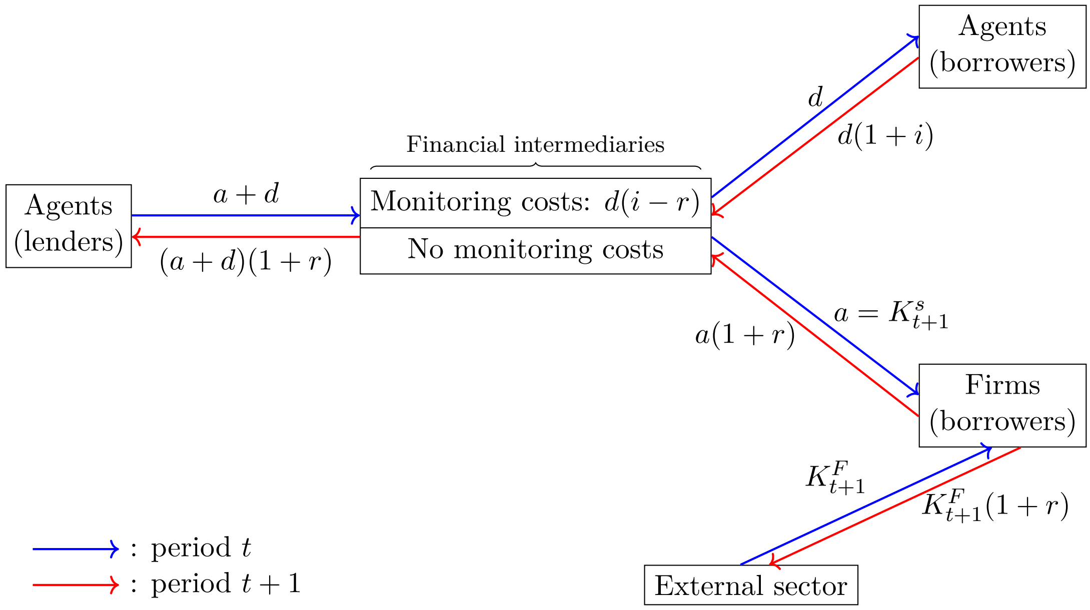

Figure 1.3: Functioning of the financial market with frictions

Table 1.1: Budget constraint of agents

Decision/Bequest	$x_t^y < h$	$h \leq x_t^y$
Not invest	**NL**	**NL**
Invest	**IB**	**IL**

types of budget constraints: **NL**, **IB**, and **IL**. The three cases are shown in Table 1. Naturally, an agent that invests in her education will borrow only if the bequest she receives (x_t^y) is lower than the cost of education (h). If the bequest is higher, she will lend the remainder. This means the comparison of x_t^y and h separate these cases. Note that the budget constraint is the same for those agents that do not invest. These cases are studied now.

- **NL (No investment – Lender)**: In this case the agent does not invest in education, and works both periods as unskilled. Since there is no consumption in the first period, the young agents is a lender. The budget constraints of this agent when young and old are respectively:[18]

$$\begin{aligned} s_t^y &= w^n + x_t^y > 0 \\ c_{t+1}^o + x_{t+1}^o &= w^n + s_t^y(1+r) \end{aligned} \tag{1.13}$$

[18] Utility is strictly increasing in both components so the budget constraint is binding.

The joint budget constraint is:

$$\frac{c^o_{t+1}}{1+r} = (w^n + x^{y,t}_t) + \frac{w^n - x^o_{t+1}}{1+r} \tag{1.14}$$

The agent maximises lifetime utility with respect to this inter-temporal budget constraint by choosing consumption and bequest when old, i.e. c^o_{t+1} and x^o_{t+1}. The first order conditions (FOCs) from this optimisation are:

$$\begin{aligned} c^o_{t+1}: \qquad & \frac{\alpha}{c^{*o}_{t+1}} - \frac{\lambda}{1+r} = 0 \\ x^o_{t+1}: \qquad & \frac{1-\alpha}{x^{*o}_{t+1}} - \frac{\lambda}{1+r} = 0 \end{aligned} \tag{1.15}$$

Note that FOCs are sufficient to guarantee a maximum because $U(\cdot)$ is concave in its two arguments. The two FOCs combined results in the intra-temporal Euler equation:

$$x^{*o}_{t+1} = \frac{1-\alpha}{\alpha} c^{*o}_{t+1} \tag{1.16}$$

Plugging this into the joint budget constraint leads to the optimal levels of consumption and bequests:

$$c^{*o}_{t+1} = \alpha \left[\left(w^n + x^{y,t}_t \right) (1+r) + w^n \right] \tag{1.17}$$

$$x^{*o}_{t+1} = (1-\alpha) \left[\left(w^n + x^{y,t}_t \right) (1+r) + w^n \right] \tag{1.18}$$

where the term in square brackets is equivalent to the total wealth of the old agent in $t+1$. This is, the agent allocates her lifetime wealth between consumption and bequest depending on their relative weight in her preferences. Equation (1.18) is of particular relevance. It indicates that the bequest left by an old agent is a function of the bequest received when young. Mathematically, this is a first order difference equation for bequests, which will be central to understand the dynamics of bequests over time later.

Finally, from the solution to consumption and bequest we can find the optimal utility of the old agent:

$$U^{NL} = \ln \left[\left(w^n + x^{y,t}_t \right) (1+r) + w^n \right] + \epsilon \tag{1.19}$$

where $\epsilon = \alpha \ln(\alpha) + (1-\alpha)\ln(1-\alpha)$.

- **IB (Investment – Borrower)**: here the agent is not rich enough to finance her education with inheritance alone, needing to borrow some resources, for which she pays an interest rate of i. She works as skilled in the second period. Each period's budget constraint for this agent is:

$$\begin{aligned} s_t^y &= x_t^y - h < 0 \\ c_{t+1}^o + x_{t+1}^o &= w_{t+1}^s + s_t^y(1+i) \end{aligned} \tag{1.20}$$

Notice the agents' decision is forward looking because it depends on the *future* skilled wage. This property will be very important when solving the skilled labour market.

The optimisation problems is solved just as above. Her optimal consumption, bequest and utility are respectively:

$$c^{*o}_{t+1} = \alpha \left[\left(x_t^{y,t} - h \right)(1+i) + w_{t+1}^s \right] \tag{1.21}$$

$$x^{*o}_{t+1} = (1-\alpha) \left[\left(x_t^{y,t} - h \right)(1+i) + w_{t+1}^s \right] \tag{1.22}$$

$$U^{IB} = \ln \left[\left(x_t^{y,t} - h \right)(1+i) + w_{t+1}^s \right] + \epsilon \tag{1.23}$$

- **IL (Investment – Lender)**: here the agent invests in education with her own resources, lending the rest to the financial intermediaries. She only works in the second period as skilled. The budget constraints for each period of her life is:

$$\begin{aligned} s_t^y &= x_t^y - h > 0 \\ c_{t+1}^o + x_{t+1}^o &= w_{t+1}^s + s_t^y(1+r) \end{aligned} \tag{1.24}$$

Again, the agent optimises in period t by taking into consideration the skilled wage of $t+1$. Her optimal consumption, bequest and utility are respectively:

$$c^{*o}_{t+1} = \alpha \left[\left(x_t^{y,t} - h \right)(1+r) + w_{t+1}^s \right] \tag{1.25}$$

$$x^{*o}_{t+1} = (1-\alpha) \left[\left(x_t^{y,t} - h \right)(1+r) + w_{t+1}^s \right] \tag{1.26}$$

$$U^{IL} = \ln \left[\left(x_t^{y,t} - h \right)(1+r) + w_{t+1}^s \right] + \epsilon \tag{1.27}$$

Optimal decisions

Having characterised optimal consumption, bequest and utility for the three possibilities, we can evaluate which decision agents take when young. As Table 1.1 indicates, young agents with $x_t^y < h$ compare **IB** versus **NL** while those with $h \leq x_t^y$ compare **IL** versus **NL**. The option that gives each young agent the higher lifetime utility corresponds to the optimal decision made by her. Old agents simply behave according to their decisions when young.[19]

- Agents with $x_t^y < h$: these relatively poor agents compares their utility when not investing (**NL**) versus borrowing to invest in education (**IB**). Thus, they invest *iff* $U^{NL} \leq U^{IB}$. This is:

$$\frac{1}{i-r}\left[w^n(2+r)+h(1+i)-w_{t+1}^s\right] \leq x_t^y \tag{1.28}$$

  The left-hand side – to be denoted as f_t – represents the cut-off point which separates those that invest from those that do not invest. In other words:

  - those that do not invest: $0 < x_t^y < f_t$
  - those that do invest: $f_t \leq x_t^y < h$

  Note that equation (2.8) assumes that, if indifferent, agents do invest. Note also that the cut-off point is **time-varying** because it is a function of the skilled wage. This is a crucial difference with respect to the basic model in GZ93, where the skilled wage is fixed (and so is f_t). Also note that f_t depends on the skilled wage of the **next period**. This comes directly from the forward looking optimisation of young agents, whose incentives to invest in education are directly related to the skilled wage they would receive when old. The consequences of these properties of f_t are studied later.

- Agents with $h \leq x_t^y$: these relatively rich agents compare their utility levels when not investing (**NL**) versus investing with their own resources (**IL**). These agents invest *iff* $U^{NL} \leq U^{IL}$. This is:

$$w^n(2+r)+h(1+r) \leq w_{t+1}^s \tag{1.29}$$

  In words, these agents invest in education <u>in t</u> whenever its reward <u>in $t+1$</u> – w_{t+1}^s – is larger than the foregone unskilled wage in both periods plus the capitalised income

[19] Strictly speaking, old agents also re-optimise when old, but because agents' preferences are time-consistent, their optimal decisions coincide with the ones made when young. The proof is straightforward and not included.

Table 1.2: Optimal investment decision for agents, given $C1$

Decision/Bequest	$x_t^y < f_t$	$f_t \leq x_t^y < h$	$h \leq x_t^y$
Not invest	**NL**	–	–
Invest	–	**IB**	**IL**

from the education spending. Unlike the previous group, here **all** agents behave equally, as x_t^y does not appear in equation (2.10). Let us denote the left-hand side of equation (2.10) as $C1_{min}$. This inequality means **all** agents in this group will invest in education in t whenever $w_{t+1}^s \geq C1_{min}$. Since it is reasonable to expect that rich agents invest in human capital, this inequality is a condition imposed on the model, to be denoted **Condition 1**, $C1$.[20] Note that *Condition 1* is dynamic and forward-looking, since it is based on w_{t+1}^s.

Table 1.2 summarises the above results. This categorization is used throughout the rest of the chapter.

Poorest agents as lenders?

The framework developed so far implies that the poorest agents are lenders (**NL**), allowing them to build up a bit of extra wealth through the financial sector. How unreasonable is this feature of the model? In light of the historical evidence, perhaps not so much. In effect, whilst in the early stages of capitalistic development private banking was not accessible to the poor, the creation of the postal savings system (in which savings account were offered in post offices) in the UK (1861) and in the US (1911) was aimed directly to citizens without access to banking, including the poorest individuals. Such accounts, offering between 2% and 2.5% of interest per year, came to accumulate quite a significant amount of funds ((United States and Congress, 1910); (Jaremski, Perlman and Schuster, 2016)). Although it is not clear to which extent

[20]Considering that a proportion of relatively poor young agents are borrowing in order to invest in education (those with $f_t \leq x_t^y$), this assumption makes perfect sense within the model. Moreover, there is a convergence consideration. Imagine a dynasty with a young agent receiving a bequest just below the education cost h. She will invest in education and the next period she will leave a bequest higher than h. But if equation (2.10) does not hold the new young agent will not invest, then it is likely she leaves a bequest below h (if not her, then her children or grandchildren, etc). Under a plausible calibration of the model this cycle will continue forever with dynasties switching endlessly from skilled to unskilled and vice versa. This economy never converges.

these accounts were used by the poorest in society, it certainly expanded basic saving instruments to financially disenfranchised individuals.

Having said this, how consequential for the model is that the poor can lend? Put differently, how would the dynamic of the variables and the steady state change if we were to impose, for example, no access to a saving account by those who do not invest in education? It is trivial to show that the bequest left by such individual is $x^o = (1-\alpha)(2w^n + x^y)$. In the long-run, these bequests converge toward $x^n = \frac{1-\alpha)}{\alpha} 2w^n$ (which is the point where $x^o = x^y$). However, if at some point along the development process enough number of individuals invest in education, and the externality is high enough to sustain a relatively high skill premium, f_t will be below x^n, which means that even the poorest agents will invest in education by borrowing. The long run steady state is the same, even though it might take longer to achieve.

The above solution assumes that poor individuals have access to the financial market as **borrowers** (but not as savers), regardless of the collateral these individuals have (in terms of inherited wealth). A more realistic specification of the model would assume complete exclusion of the financial markets by those without wealth above a given threshold $\hat{x}$. For instance, $\hat{x}$ could be defined as a proportion of the unskilled wage (e.g. $\hat{x} = \delta w_n$). The consequences of such specification vary. On the one hand, if $x^n > \hat{x}$, all these poor dynasties will at some point transit toward the skilled sector, as f_t falls. This case is no different with respect to the original specification of the model, except for a slower convergence towards the steady state of full development. On the other hand, if $x^n < \hat{x}$, all those dynasties which initial ($t = 0$) bequest is below $\hat{x}$ will always be part of the unskilled sector, ending up with a steady state bequest level of x^n. Importantly, this is the case regardless of how low f_t gets, as the financial constraint overrides any young agent's desire to invest. In other words, even if young agents would like to become skilled, their collateral is not high enough for them to access the financial market. The steady state is quite different now, as not every dynasty is composed of skilled old workers; a section of the population (given by the initial bequest distribution) remains unskilled.

Before concluding, some historical remarks are of relevance. First of all, there is good evidence that among early European and US universities, financial aid to poor students, usually from wealthy philanthropist was widespread ((Fuller, 2014)). Naturally, there were only very few the poor children lucky enough to go to these institutions. Grand-scale schemes of financial support through government-backed loans, for example, are much more recent (in the US, they came up in 1958). Thus, whilst on the one hand credit restrictions are likely to be of historical importance, financial development together with government provision of loans have to some extent lowered such restrictions. This can be though as a fall in δ over time, with $\lim_{t\to\infty} \delta = 0$. In other words, the simplification used in this chapter of $\delta = 0$ ($\hat{x} = 0$) can be thought as representing a very long steady state, which even today has not being achieved in any country. But ultimately, there is no much difference between this simple setting and a more realistic where δ (and $\hat{x}$) falls as the economy grows.

A note on the fixed cost of education

Naturally, in a model where the skilled premium evolves endogenously, assuming that the cost of education is fixed is quite unrealistic. However, it is not hard to see that this assumptions is much more benign that what it appears. To see this, assume that the cost of education is a fixed proportion λ of the skilled wage individuals expect when working.[21] This is, $h(w^s_{t+1}) = \lambda w^s_{t+1}$. The new definition of f_t is:

$$f_t \equiv \frac{1}{i-r}\left[w^n(2+r) - w^s_{t+1}\left[1-\lambda(1+i)\right]\right] \tag{1.30}$$

Whilst in the case of a fixed h the derivative of f_t with respect to w^s_{t+1} is negative (of crucial importance in the development process, as shown later), in the case of endogenous education cost, such derivative depends particularly on the value of λ. In effect, three cases exist:

$$\frac{\partial f_t}{\partial w^s_{t+1}} \begin{cases} = 0 & \text{if} \quad 1 = \lambda(1+i) \\ < 0 & \text{if} \quad 1 > \lambda(1+i) \\ > 0 & \text{if} \quad 1 < \lambda(1+i) \end{cases} \tag{1.31}$$

[21] Using the current skilled wage introduces a second source of dynamics into the model – the first being bequests – which obscures considerably the development and solution of the model. It is not unreasonable to think that education costs partly reflect the *expected* return of education, and not only current returns.

The first case is equivalent to that one in GZ93, where f_t is fixed. Here, the economy can settle in a steady state with some technological specialisation, where some workers remain unskilled. The second case is roughly equivalent to the one in this chapter, except that the derivative is *smaller* (compare equations 2.8 and 1.30). This means a slower response of f_t to changes in the skilled wage, which in turn means either a slower transition to complete development or that full specialisation requires more stringent conditions in terms of bequests and degree of externality. Finally, the third case is that where an increase in the skilled wage leads to **less** agents becoming skilled (recall that young agents with $x_t^y \geq f_t$ invest in education). This is counterintuitive, and might lead to oscillations in L_t^s.

It might seem reasonable by now that the second case is the most interesting. In effect, such case requires $\lambda < \frac{1}{1+i}$. For an **annual** interest rate on borrowing of 10% (as used in the simulations), the **annual** education cost must then be below 90% of the **annual** salary of future skilled workers. This is by no means an unreasonable assumption.[22]

Thus, if we were to relax our assumption of fixed education cost, and replace it by the most reasonable alternative, which is the second case in equation (1.31), the conclusions do not vary very much. In effect, for a sufficiently high externality or bequest levels, the same convergence toward a fully develop society (and a Kuznets curve) can be achieved. It is true that the transition process might be slower under endogenous h than under fixed h, *ceteris paribus*, but given the rather arbitrary nature of time periods in the OLG structure, the speed of transition is not a very insightful dimension anyway. Therefore, treating h as exogenous is a convenient simplification of the model, without reducing significantly the insights that arise from it. In other words, the central results do not rest on such assumption.

1.3.5 Supply of skilled labour, unskilled labour and capital

As highlighted above, the level of bequest fully characterises agents' educational choices. This enables to find the supply of skilled workers in any period. At the same time. the

[22]Notice that **Condition 1**, which is required for rich inheritors to invest in education (and for the model to be interesting), puts an upper limit to λ. It is trivial to show that this boundary is $\lambda < \frac{1}{1+r}$, which still gives a cost of annual education below 100% of annual skilled wage. Thus, the margin of manoeuvre for different cases is quite reduced. It does not seem less arbitrary to assume λ between 90% and 100% than below 90%.

supply on unskilled labour is defined. Since savings define the source of local funds, the supply of capital is also characterised.

Supply of skilled labour

As Table 1.2 shows, the cut-off point f_t separates those young agents that do invest in education from those that do not invest. Thus, in order to determine the supply of skill workers in $t+1$, it is enough to look at the number of young agents with inheritance above f_t in t. **This inter-temporal relationship is crucial**. In effect, it is in period t that agents choose whether to become skilled or not. If they do, their are defining the supply of skilled labour in period $t+1$. The amount of skilled workers in period t was already decided in $t-1$, when the now skilled population decided to invest in education. This inter-temporal property of the supply of skilled labour means we need to distinguish between the *ex-ante* and the *ex-post* supply. The former is based on the prospect of what the supply will be in the next period, according to the educational decisions of agents. This supply varies with the skilled wage. The latter is equivalent to the number of agents who are effectively skilled in a period, and it is inelastic. These two are analysed below.

- *Ex-ante* supply:

Recall that f_t is a function of the skilled wage in $t+1$, as equation (2.8) shows. This is, $f_t(w^s_{t+1})$. In consequence, the supply for skilled labour in period $t+1$ is:

Ex-ante supply of skilled labour

$$L^s_{t+1} = \mathrm{F}_1(w^s_{t+1}) \tag{1.32}$$

where $\mathrm{F}_1(\cdot)$ denotes an arbitrary function. Unfortunately, this function cannot be written algebraically because it depends on the distribution of bequests received by the young agents in period t, which varies in arbitrary ways. In any case, this supply has some stable properties, shown in Panel (a) of Figure 1.4. In particular:

- The point $(0,0)$ is part of the curve, because if the wage is zero, no agent would invest in education. This is the only exception to **Condition 1** allowed.

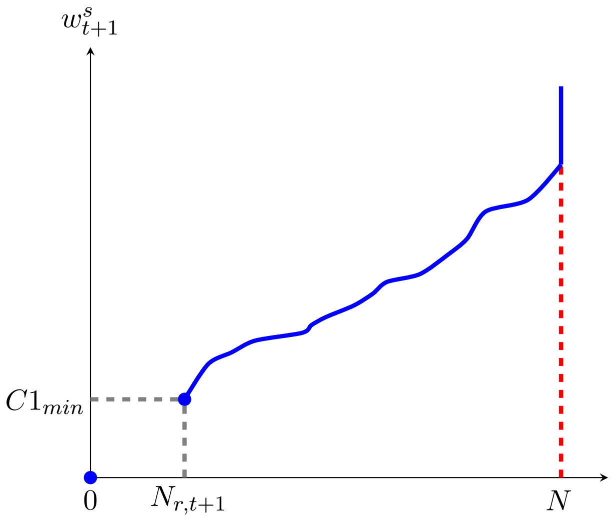

(a) *Ex-ante* supply of skilled labour

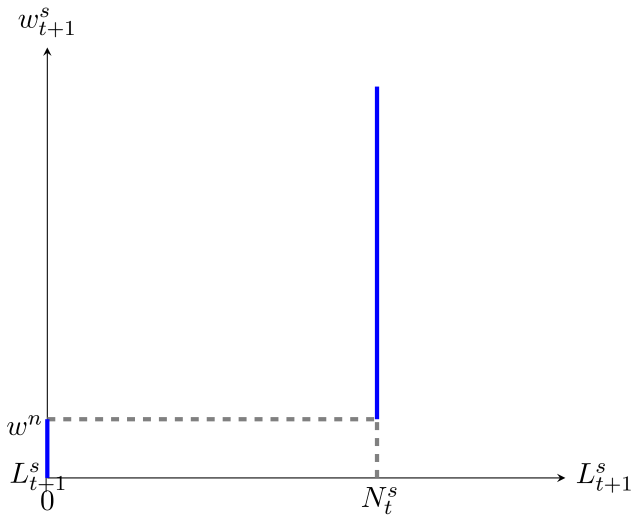

(b) *Ex-post* supply of skilled labour

Figure 1.4: Example of the supply of skilled labour in period $t + 1$

- Since **Condition 1** holds for any positive labour, the skilled wage cannot be lower than $C1_{min}$. At this level, all those with $x^y_t > h$ (the richest) invest in education. Thus, the supply starts at $N_{r,t+1}$, where $N_{r,t} = \sum_{j \in N} \text{I}(x^{y,j}_t > h)$ and $\text{I}(\cdot)$ is the indicator function.[23]

- For the remaining of the young agents (those with $x^y_t < h$) the supply is increasing in the wage because the cut-off point at which they invest, f_t, is decreasing in w^s_{t+1}. In other words, for a higher w^s_{t+1}, more agents invest in education, which results in an upward slopping curve. The shape of the curve in this section depends on the distribution of bequests among these agents in period t (B_t). The more concentrated (disperse) bequests are, the more horizontal (vertical) the supply is. For instance, if several agents inheriting an equal bequest, the supply is horizontal along that section.

- At N (the generation size) this supply becomes perfectly inelastic.

Notice that the supply curve represents an ordering of agents according to their inheritances. For a very low skilled wage only the richest will invest in education. As the

[23]Strictly speaking, any point along the dashed line is not part of the supply because at that wage **all** rich invest in education. To introduce this section into the supply requires, for example, to add a constraint in the number of places available in educational institutions, where selection could be done through inherited wealth. In turn, only the richest among the rich fill up educational places, with a group of other relatively rich agents remaining unskilled against their will.

wage increases those agents that follow on the inheritances distribution will invest, until the wage is so high that even the poorest agents invest in education. Additionally, the curvature of this supply can be any combination of convex and concave sections, albeit this supply is never decreasing. The curvature depends on the dispersion and level of bequests. A convex (i.e. positive second derivative) occurs when we move from relatively disperse level of bequest toward closer ones, whereas a concave one (i.e. negative second derivative) occurs when bequests are becoming increasingly more disperse, this as we move along the bequest distribution. Finally $N_{r,t+1}$ and the upward slopping section of the supply might change over time, as both w^s_{t+1} and the distribution of bequests evolve. More on this later.

- *Ex-post* supply:

Once some young agents decide to become skilled, they spend their youth in education. Let us denote the number of agents who decided to become educated in period t as N^s_t. The next period, they enter the labour market, as qualified workers. But, will they actually join firms in the modern, skilled sector? The answer is trivial. They will do so as long as $w^s_{t+1} > w_n$. In turn, the *ex-post* supply of skilled labour has the following properties, also presented in Panel (b) of Figure 1.4:

- It is zero for $w^s_{t+1} < w_n$. This corresponds to a perfectly inelastic section.
- It is equal to the *ex-ante* supply (equivalent to the number of agents who invested in education in the previous period, denoted N^s_t) for $w^s_{t+1} > w_n$. This also corresponds to a perfectly inelastic section.
- It is undefined at $w^s_{t+1} = w_n$, between $L_s = 0$ and N^s_t. This is because agents are indifferent between sectors. Alternatively, we could assume this is an horizontal section, where individuals are randomly allocated into sectors if the equilibrium goes through this portion.

This supply can be summarised as follows:

Ex-post supply of skilled labour

$$L^s_{t+1} = \mathrm{F}_2(w^s_{t+1}) = \begin{cases} N^s_t, & \text{if } w^s_{t+1} > w_n \\ N^s_t > L^s_{t+1} > 0, & \text{if } w^s_{t+1} = w_n \\ 0, & \text{if } w^s_{t+1} < w_n \end{cases} \tag{1.33}$$

where $\mathrm{F}_2(\cdot)$ denotes an arbitrary function.

Supply of unskilled labour

The supply of unskilled labour in t depends on agents' educational decisions. In particular, it depends on how many agents decide **not** to invest in education in t. Therefore, the supply of unskilled labour can be written as a function of the skilled labour.

First, the following inter-temporal identities relating productive sectors and generations of agents hold every period:[24]

$$\begin{aligned} N &\equiv L^{y,n}_t + L^s_{t+1} \\ N &\equiv L^{o,n}_t + L^s_t \\ L^n_t &\equiv L^{y,n}_t + L^{o,n}_t \end{aligned} \tag{1.34}$$

From these it is possible to derive:

$$L^n_t = 2N - (L^s_t + L^s_{t+1}) \tag{1.35}$$

Although skilled workers in t (L^s_t) is predetermined in t (the investment decision was made in $t-1$), L^s_{t+1} is endogenous in t. Therefore, the supply of unskilled labour can be written as:

Supply of unskilled labour

$$L^n_t = \mathrm{F}_3(L^s_{t+1}) \tag{1.36}$$

where $\mathrm{F}_3(\cdot)$ is a known function (from equation 1.35).

[24] The first identity in equation (1.34) splits the young generation between those that are working as unskilled and those that are in education (to be skilled the next period). The second identity splits the old generation between those that did not invest and those that did invest in education in $t-1$. The third identity states that at every period those working in the unskilled sector are the sum of the young and the old unskilled.

Supply of capital

The supply of capital has two sources: domestic and foreign. As in a standard OLG model, the supply of domestic capital for period $t+1$ – K^s_{t+1} – comes from the savings of the young agents in t (as Figure 1.3 shows). More precisely, old agents own the capital stock (and firms) in period t. That period, young agents aggregate savings are defined. At the same time, those savings are used to buy old agents' capital, income which the latter use to consume and leave bequest. Notice that this sell transfer the ownership of capital to the young generation, equivalent to the old one the period after, restarting the cycle. If young agents savings exceed capital available, then they use those savings to invest in new capital. The standard dynamic condition for capital applies:[25]

$$K^s_{t+1} = K^s_t(1-\delta) + I_t$$

where K^s_{t+1} is the sum of young agents' savings. More precisely:

Supply of domestic capital

$$K^s_{t+1} = \sum_{j=1}^{N} s^{*y}_t \quad (1.37)$$

In this economy open to the international capital market the supply of foreign capital – K^F_{t+1} – is just the residual of any imbalance between the domestic demand and supply of capital. This is:

$$K^F_{t+1} \equiv K^D_{t+1} - K^s_{t+1} \quad (1.38)$$

Appendix A shows the rest of the macroeconomic identities, and shows that any persistent imbalance in the Balance of Payments (due to a persistent imbalance in the capital account reflecting capital flows) is sustainable even in the long-run. Thus, the model is internally consistent.

[25] Notice that complete access to capital markets breaks the dependence between periods that the dynamic of capital usually introduces (e.g. in Solow or Ramsey models). This is, regardless of wheter capital fully depreciates every period or not, any excess or shortage of supply with respect to demand is supplemented by international capital flows.

Table 1.3: Aggregate demand and supply for each factor

	Aggregate demand (by firms)	Equation	Aggregate supply (by agents)	Equation
Unskilled labour	$w_t^n = w^n$	(1.2)	$L_t^n = \mathrm{F}_3(L_{t+1}^s)$	(1.35)
Skilled labour	$w_t^s = A^s \left(L_t^s\right)^{\psi-\theta} (1-\theta) \, K_t^{\theta}$ (partial equilibrium)	(1.9)	$w_{t+1}^s = \mathrm{F}_1^{-1}(L_{t+1}^s)$ (ex-ante)	(1.32)
	$w_t^s = A^s \left(L_t^s\right)^{\frac{\psi}{1-\theta}} (1-\theta) \left(\frac{A^s\theta}{r}\right)^{\frac{\theta}{1-\theta}}$ (general equilibrium)	(1.10)	$w_{t+1}^s = \mathrm{F}_2^{-1}(L_{t+1}^s)$ (ex-post)	(1.33)
Capital	$k_t^{D*} = \left(\frac{A^s \left(L_t^s\right)^{\psi}\theta}{r}\right)^{\frac{1}{1-\theta}}$	(1.8)	$K_{t+1}^s = \sum_{j=1}^N s^*{}_t^y$ (domestic)	(1.37)
			$K_{t+1}^F \equiv K_{t+1}^D - K_{t+1}^s$ (foreign)	(1.38)

1.3.6 Summary of equations

The demand and supply for each factor has been fully characterised. Table 1.3 summarises the results. These equations are used in the next section to find the equilibrium in each factor market.

1.4 Equilibrium in the factors' markets

Recall that in this model there are four markets: the market for capital, for skilled and unskilled labour, and the final goods market. The financial market is trivial, due to the complete access to international capital markets. In other words, it will always be in equilibrium. The final goods market is a "residual" market from the perspective of the Walras Law. This is, once we solve the equilibrium in the two labour markets, the goods market is solved too. Consequently, to resolve the model, it is enough to focus on the labour markets. Still, ss the previous section indicated, the two labour markets are intrinsically connected, due to the educational choice young agents make. This means that we can focus solely on one market, which is that for skilled labour. Since the skilled wage is endogenous, this market is where most of the action is happening. In turn, this section will focus first on finding the equilibrium in the skilled labour market. Then, it shows how the rest of the endogenous variables can be derived from it.

One particularity of the skilled labour market is that both the supply and the aggregate demand for skilled labour are upward slopping. In effect, the number of agents that would invest in education is a direct function of the skilled wage expected for the next period. Also, firms in this sector pay a wage which is an increasing function of the number of skilled workers available that period due to the externality (recall the upward slopping nature of the "total equilibrium" demand for skilled labour). This might introduce multiple equilibria in this market, and multiplicity of solutions to the model. Nevertheless, it is shown that there is always one equilibrium which is superior to the rest (in the Pareto sense), meaning the solution to the skilled market, and to the rest of the model is actually unique.

Another particularity of this market is the existence of an *ex-ante* and *ex-post* supply of skilled labour. This however has no practical consequences, because of perfect competition. Still, they are important in defining the equilibrium of the model.

1.4.1 The skilled labour market equilibrium

The analysis that follows focuses on the *ex-ante* supply of skilled labour. This is, the goal is to characterise **how many agents would invest in education in period** t. The previous section showed that this decision depends crucially on the skilled wage, thereby defining the *ex-ante* supply of skilled labour.

The equilibrium in this market for a given period is found by the combination of the supply and demand for skilled labour. Is there always an equilibrium in this market? Is it unique? The answer lies in the curvature of the "total equilibrium" demand for skilled labour in combination with the shape of the supply curve. Here it is assumed that $\psi < (1-\theta)$, making this demand increasing and strictly concave.[26] This condition rules out externalities being too great, allowing for a smoother development transition process. If not, the model can easily turn uninteresting, as all young agents would

[26] Appendix B explores in detail this demand, showing that it is always increasing, and its concavity depends on the sign of $\psi - (1-\theta)$.

become skilled in the first period. Importantly, this assumption does not drive the multiplicity of equilibria found later. In fact, multiple equilibria might arise for any parametrisation, as long as $\psi > 0$, i.e. the externality exists.

Panels (a) and (b) in Figure 1.5 plot in a single graph both the *ex-ante* supply introduced in Figure 1.4 and the "total" demand for skilled labour, for two different economies (or stages of the same economy). Panel (a) represents a "middle income" economy in the sense that many but not all could become educated (point C). In contrast, Panel (b) represents a "rich" economy where all become educated (point D), either because young agents inherit a high bequest, the skill premium is high enough to motivate all to become educated, or both. Notice that the supply curve is "lower" in Panel (b), reflecting the fact that agents receive larger bequests. This is, to produce the same number of skilled workers, the "rich" economy requires a smaller skill premium. In both cases, there is one equilibrium with positive skilled labour, and a trivial equilibrium without skilled workers.[27] Regarding the latter, if no agent is skilled, firms in the modern sector do not produce (due to the Cobb-Douglas technology), and the skilled wage is zero. Demand and supply coincide. Regarding the equilibria with positive employment in the skilled sector, Point B represents an unconstrained equilibrium in the sense that there is a section of the young generation not investing in education. Conversely, point C represents a constrained equilibrium in the sense that even by adding poor agents (e.g. equal to the poorest in the economy), they would become skilled. Clearly, there always exist a trivial equilibrium (point A), solving the existence problem. In fact, a non-trivial equilibrium might not exist. This is the case if the externality if relatively low and the economy is relatively poor, as shown in Panel (a) in Figure 1.6. The supply is the same as that in Panel (a) of Figure 1.5, but the demand is less steep, because ψ is lower. This contrasts with the case of an initially rich economy, where even a low externality might still yield a non-trivial equilibrium in a rich economy. Panel (b) of Figure 1.6 shows this example. It reproduces the aforementioned supply of the "rich" economy, while only changing the demand. As said, the richer the economy is, the flatter the supply of skilled labour is because more young agents receive an inheritance

[27] Recall that the point where the demand crosses the "hypothetical" horizontal section of the supply is not an equilibrium in the strict sense. It requires to introduce a restriction in the number of educational places available. Notice however that removing this restriction is a complete win for all parties in the economy. This is studied later.

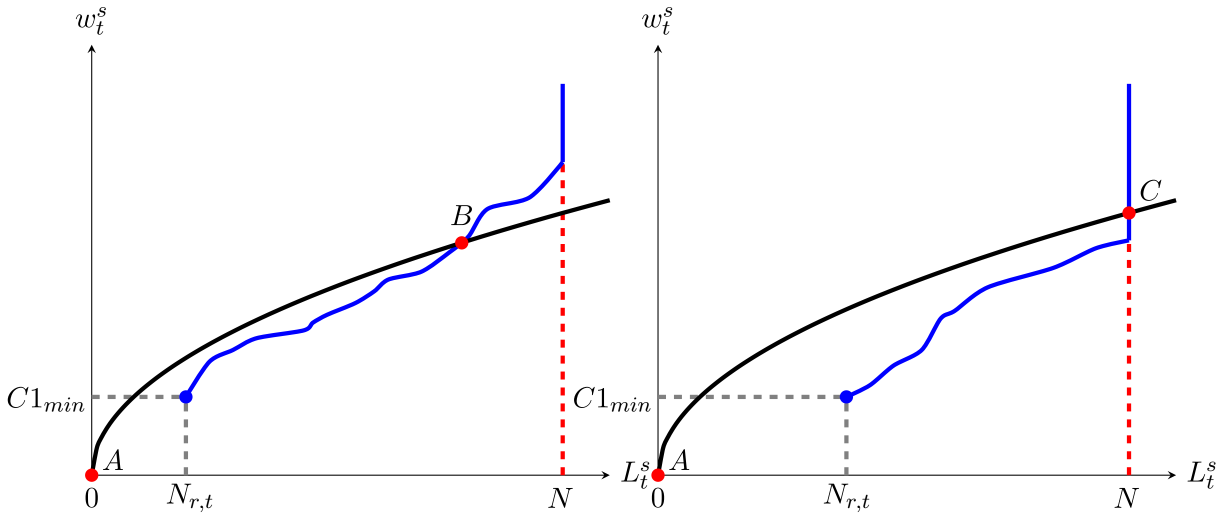

(a) Middle income economy (b) High income economy

Figure 1.5: Equilibria in the skilled labour market for two different supplies, and a relatively high level of externality

above h (i.e. the horizontal region of this supply is long) and very few are far away from this threshold (i.e. the extent of vertical increases is small). In turn, even for a relatively low externality, firms are able to pay a wage (rather low) that motivates some rich enough agents to become skilled.

This provides interesting insights. In order to produce with modern technology, rich countries have an advantage because agents can afford to acquire the skilled needed for it, even if the pay-off in terms of higher wage is not too great. On the contrary, poor countries require a higher incentive in terms of wage, which might be unprofitable for local firms. Thus poor countries might only introduce modern technology in sectors that have high externalities.

Multiple equilibria

It is also possible for the model to produce multiple non-trivial equilibria. There is no simple rule that defines when this happens, but some insights arise from the analysis of the curvature of the supply and demand curves. First of all, having assumed a strictly concave "total" demand (i.e. $\psi < (1-\theta)$), it is very unlikely for a convex supply curve to produce multiplicity of equilibrium. A convex supply comes from an economy with a relatively large middle class (in terms of those with bequests within a certain range),

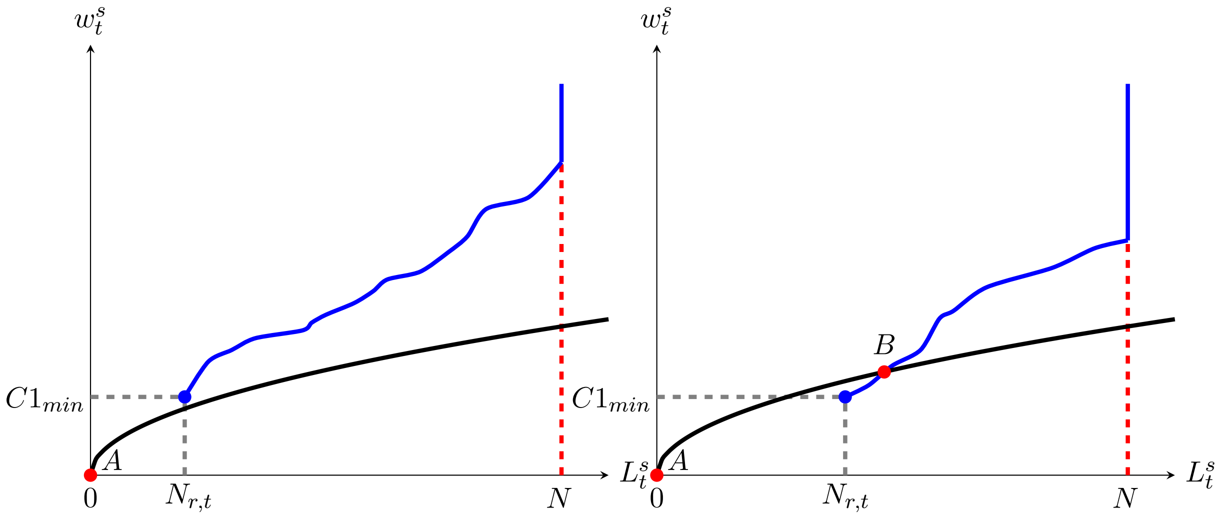

(a) Middle income economy

(b) High income economy

Figure 1.6: Equilibria in the skilled labour market for two different supplies, and a relatively low level of externality

and relatively few poor agents. This is shown in Panel (a) of Figure 1.7. Here, for a small increase in the skilled wage, a large group of agents would choose to become skilled. Still, it is clear from the graph that a small number of rich agents (i.e. small $N_{r,t}$) would introduce another equilibrium to the market.

Multiplicity of equilibrium might be less likely too in an economy with a small middle class, where the majority of agents are poor. In this case, the supply curve is concave, as Panel (b) of Figure 1.7 shows. Here, a large skilled wage is needed to induce young agents to borrow to invest in education. Although not shown, it is clear that multiplicity of equilibria is more likely to happen in an economy where the supply has several convex and concave paths. This is an economy where the distribution of bequests is less smooth; for example, if bequest are densely concentrated at several values, distribution which *pdf* has several spikes. How likely this wealth distribution is in practice is an empirical matter, but it might not be the most common shape. All in all, the bottom line of this analysis is that smooth distributions might not lead to multiplicity of equilibria, as in Figure 1.7.

Now, if an economy faces multiple equilibria, which one prevails? This depends on the decisions (whether to invest in education or not) of every agent. Yet, as every agent's choice affects the pay-off structure of every other agent, the issue becomes essentially a

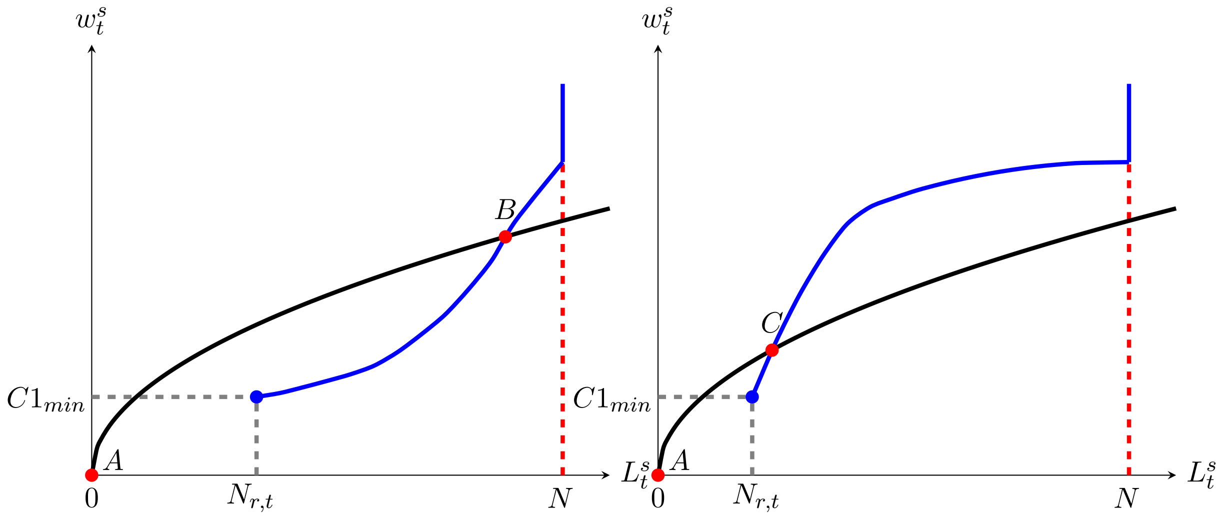

(a) Economy with a large middle class

(b) Economy with a small middle class

Figure 1.7: Equilibria in the skilled labour market for two different economies

game theory problem. Such game involve multiple players and two actions (to invest or not), and it is a complete information (i.e. agents now all payoffs), asymmetric, non zero-sum, simultaneous game.

The approach to solve this game is to categorise workers into groups. To do this, let us assume, without loss of generality, that there are two (non-trivial) equilibria, derived from the intersection of demand and supply for skilled workers: one with low L^s_{t+1} and one with high L^s_{t+1}. In this setting, the population of young agents can be split into three groups: (i) those who remain unskilled in both equilibria (the poorest), (ii) those who choose to become skilled anyway (the richest), and (iii) those who in one equilibrium would prefer to be educated, whereas in the other would not. The exercise that follows aims to establish whether the strategic interaction of agents leads to one equilibria to be always chosen. Put it differently, we can see if, starting from each equilibrium, each type of agent would be better off by changing their decisions, given what other agents choose. If one equilibrium is always chosen (i.e. is Pareto superior), we conclude that such equilibrium prevails.

Since simultaneous games with three (type of) players are graphically cumbersome, we can reduce the game to two (type of) players by noticing that for group (i) (those who never invest), the action of not investing in education is a dominant strategy. In effect, for either the low or high equilibrium, these agents are too poor to invest in education (borrowing cost is too high). In consequence, **regardless of what any other type of player does**, they are better off by remaining unskilled – the dominant strategy.

Having only two players, we can present the game as follows:

		Group (ii)	
		Not invest	Invest
Group (iii)	Not invest	(a, b)	$(a, \underline{c})$
	Invest	(d, b)	$(\underline{e}, \underline{f})$

Group (ii) refers to those who in both equilibria would choose to invest, and group (iii) refers to those who would change their decisions between the two equilibria. The pay-offs are given in the matrix. First, notice that the pay-off of those who do not invest (a and b) is independent of other agents decisions (like for group i). Second, because a higher number of skilled workers mean a higher skilled wage, it is true that $c < f$ and $d < e$. This is, the pay-off of a group that invests in education increases if the other group also invests in education. Third, the low equilibrium (where only group (ii) invest in education) is represented by the pay-offs (a, c), whereas the high equilibrium is represented by the pay-offs (e, f). Fourth, by definition of each group, it must be true that $b < c < f$ (i.e. those in group (ii) invest regardless of what the other do), and that $a < e$ (i.e. if group (ii) invests, group (iii) also invest). The only information we do not know is whether group (iii) is better of investing when group (ii) is not investing (the first column of the game, or a *versus* d). However, this is not a possible scenario. As the simple comparison of pay-offs show, the only pure strategy equilibrium is that where both groups of agents invest in education. Actually, this is unsurprising. **Condition 1** means that invest in education is a dominant strategy for group (ii). Therefore, it follows that for group (iii), invest in education is also preferred.

In consequence, the high equilibrium (Invest, Invest) is a *dominant-strategy, pure-strategy, Nash equilibrium.* This is, there is not really a strategic component in agents interaction. As long as they are rational agents, group (i) will choose not to invest and the rest will choose to do so. The above proof follows through for any number of equilibria; namely, rationality ensures the one with the highest number of skilled workers is chosen.

In summary, an economy with a relatively low externality might never produce with the modern technology, unless it has a sizeable amount of wealthy inheritors willing to accept a low skilled wage. The rest finds this wage too low in order to pay back the large debt required to become educated. An higher externality is enough to kick-start the transition from traditional to modern technologies, even in poor societies. Here, some agents do borrow to finance their education, in a process that might well be the beginning of the creation of the middle class in the economy. Finally, any possible multiplicity of equilibria is inconsequential, because rational agents would always conduct the economy towards the Pareto Optimal one.

Ex-ante and ex-post equilibria in the skilled labour market

The previous section studied the decision of agents to become skilled, a decision taken in period t. Now, in the next period, there are N^s_{t+1} skilled workers. Here, the *ex-post* supply operates. The central question is, does the *ex-ante* and *ex-post* equilibrium coincide? This is, would firms pay w^s_{t+1} equivalent to the *ex-ante* level, or would they pay $w^n + \epsilon$? The latter would certainly be beneficial for firms, because they would still use the full input of workers' labour in production, paying them well below their marginal product. This however is not sustainable. In effect, if firms could "exploit" workers in such way, their demand for skilled labour would be much higher. They would like to hire workers until the marginal product of labour is $w^n + \epsilon$, with $\epsilon \approx 0$. Then, in a competitive market, firms will bid for workers by rising ϵ. Unexpectedly, the exact point where the demand for skilled labour by firms equal N^s_{t+1} is at the *ex-ante* w^s_{t+1}. This means the assumption of perfect competition makes the model dynamically consistent in the sense that the permanent fooling of agents is not possible. This saves the need to introduce any external commitment device, e.g. related to reputation.[28]

Although not shown, it is trivial to show that the *ex-post* and *ex-ante* equilibria are the same using the graphical analysis of supply and demand. The point where the demand crosses the *ex-ante* supply defines N^s_{t+1}. The *ex-post* supply is inelastic at N^s_{t+1}. Therefore, by definition, it crosses the demand at the same skilled wage level than the *ex-ante* supply.

[28]If the *ex-ante* and *ex-post* equilibria were different, agents could "learn", and ask firms to commit to hiring them the next period.

1.4.2 Solution of the full model

Once the solution to the skilled labour market is found, everything else becomes known. In particular, the educational choice defines optimal L^s_{t+1} and L^n_t, according to equation (1.35). Similarly, L^s_{t+1} solves the optimal level of capital in $t+1$, from equation (1.8) (recall the stock of capital in t is already known). The supply of capital is obtained from aggregating agents' savings (known from their educational choices, which defines their income and savings), whereas the net foreign capital flow equilibrates the latter two.

It is useful to point out the timing of the model's solution. First, notice r, i, and w^n are exogenous throughout the periods. The economy starts in $t = 0$, with the old generation's skill status given (recall that the educational decision is chosen when young). In particular, L^s_0 is known. Also, they already have decided how much to consume, and how much to bequest to the newly-born, young generation. This is, the initial distribution of bequests B_0 is known. Finally, their savings in the previous period fixed the level of capital available in $t = 0$, capital owned by the old (recall capital is defined as a stock in the beginning og the period). In consequence, from equation (1.10), the skilled wage in $t = 0$ is known. Production in the modern sector is also known.

Regarding the young, they have received an inheritance and must decide whether to invest in human capital or join the labour market as unskilled. Given the public information available for young agents regarding contemporaneous bequests, preferences, and production functions, they (like we have done here) can compute the demand and supply for skilled labour in order to find the **ex-ante** equilibrium for the next period, $t+1$. Recall that this is unique, given agents' utility maximisation behaviour. Denote this equilibrium as $\{w^{*s}_1, L^{*s}_1\}$. This predicted wage sets f^*_t, allowing each individual young agents to decide whether or not to invest in education, given her individual bequest received. Naturally, an exact number of L^{*s}_1 agents invest. As argued before, perfect competition ensures the **ex-post** and **ex-ante** equilibrium coincides.

Now that the young's vocational discernment is over, the unskilled young workers join the unskilled old ones. This is, unskilled labour for $t = 0$ is found, which then determines output in the unskilled sector. GDP and therefore consumption are resolved, finalising the analysis for $t = 0$.

Very importantly, notice that by the end of $t = 0$, the following $t = 1$ variables are known: B_1, L_1^s, w_1^s, $L_1^{o,n}$, and K_1^s. Even more, they are the same initial conditions which the model started with in $t = 0$. Thus, iterating the above procedure, the solution for period $t = 1$ can be fully resolved, together with the initial conditions for $t = 2$. Clearly, by induction, the whole future of this economy is known. This highlights the deterministic nature of the model, where the solution path of all endogenous variables relies entirely on the initial conditions.

1.5 The Kuznets curve

The previous section showed that there is always a trivial equilibrium in the skilled labour market, namely $L_t^s = 0$. Furthermore, for a high enough level of the externality, there is at least one non-trivial equilibrium (i.e. $L_t^s > 0$). This section focuses on the equilibrium in this market over time, which then characterises the development path the economy takes, in terms of both GDP and inequality. Perhaps unsurprisingly, there is no structure in the model that ensures a unique development path. For example, the skilled wage and therefore the number of skilled workers could well decrease or increase over time, or move cyclically. In effect, probably any path can be obtained via a particular parametrisation of the model. Since the goal here is to reproduce a Kuznets curve, the analysis focuses particularly on the conditions that lead to it. This should not distract from the fact that the model is flexible enough to produce alternative outcomes in terms of inequality and growth.

Without doubt this is the central part of this chapter. Here, the consequences of the adjustable skilled wage (a by-product of the externality in skills acquisition) for the short-run and long-run dynamic of the endogenous variables become clear.

1.5.1 Definition

A Kuznets curve is an inverted U-shape pattern of (generally income) inequality for different stages of development of countries: low inequality for relatively poor and rich economies, and high inequality for middle income economies. In this model, the level of development is calculated from output per capita ($\frac{Y^n+Y^s}{2N}$), whereas inequality is calculated from both income (wage) of the working population (including young and old), and wealth (bequest) of the young generation.

Since the population is fixed, GDP per capita grows as long as GDP grows. Moreover, since the productivity of the traditional sector is fixed (at w^n), the only way to expand the economy is through an employment shift toward modern technology. This shift comes from new generations investing in education. As such, development in this economy is entirely driven by workers' decisions. Recall that the externality is not captured directly by firms (they have constant returns to scale and therefore zero profits). Thus, they have no incentive to stimulate this reallocation via, e.g. a training program that produces skilled workers. This model is then an extreme version of general training, one that only benefits the worker but not the firm (Becker, 1962).[29]

The Kuznets development process

As it turns out, the process of development available in this economy, namely a transition from traditional to modern technology, can easily produce a non-linear inequality pattern over time. This is, the Kuznets curve is a rational outcome of the model. Consider a subsistence, closed economy in $t = 0$. Here, everyone is relatively poor, unskilled, and all production is carried out using traditional technology. Inequality of income and wealth are low (unless we assume some few extremely wealthy individuals, like a royal family). If agents gain utility from bequests, they might still save towards inheritance, according to equation (1.18). If not, the level of bequests are very low. In any case, there is no capital in this economy, meaning there is no demand for loans. In turn, $r = 0$.

In the next period, the economy opens, and a new, modern technology becomes available. This technology requires capital, introducing a demand for loans. Similarly, the opportunity to become skilled through education also introduces a demand for

[29] Naturally there are alternative incentives possible. For example, in a more general setting where workers are not fully rational, information is imperfect, or there are uncertainties, it is likely that many workers underinvest in their education. A government interested in development would aim to overcome this, and exploit the human capital's externality in full, either by providing (compulsory) public education (leading to the skilling of the workforce), or loans, both financed through taxation. The latter means a drag on workers' income prospects, playing against their own incentive to get skilled. Naturally, their balance depends on the particular parametrisation. This is explored in the next section.

loans; the financial market now exists. Savings are rewarded at the international rate, r. Agents have larger incentives to save now. The modern technology has the particularity of benefiting from the aggregate level of education of the population (the knowledge externality). As discussed, for an externality high enough, a group of young agents, knowing the bequest they are to inherit, decide to invest in the skills required to carry out the production. This group might include both rich lenders (e.g. the *aristocracy*) and not-so-rich borrowers (e.g. the *bourgeoisie*). They become skilled, joining technologically advanced firms in the next period. Strictly speaking, these are the kick-starters of the development process, as the number of skilled workers and output increases.

As said, to produce sustained growth, it is required that an increasing number of agents become skilled over subsequent generations. At first glance, it might seem trivial that this happens as long as the skill premium ($\frac{w^s}{w^n}$), which represents the benefit of becoming skilled, widens over time. Yet, recall the skilled wage is an outcome of the model; the *ex-post* result of an *ex-ante* analysis. Its periodic increase merely **reflects** the fact that education is becoming a convenient life choice for more young agents. The real factor leading to an expansion of L^s over time (and subsequently, of w^s) is an **ever wealthier young generation**. This is, when more parents are able to pass on more wealth than that which they received from their parents a generation ago, the economy is likely to expand. This is because the only time-varying factor defining the educational decision of agents within a dynasty is their inheritance. In fact, it is not only their own inheritance, but that of the whole generation. The wealthier other young agents in the generation are, the more likely a given agent is to become skilled, because the externality of knowledge expands the skill premium. Still, *ceteris paribus*, each individual's inheritance is the solely driver of education. To put it differently, **if the distribution of bequests is stable, the economy is in a steady state**.

As it turns out, the introduction of this new technology leads to an general increase in bequests (equivalent to a fall in the skewness of the distribution of bequests). This is evident for those young agents which invested in education, and whose parents were unskilled. The bequest they leave will be larger than the one they received (which motivate them to invest in education). This is also true for those who, like their parents, remain unskilled. Even though they earn the same (w^n), their savings are now rewarded at $r > 0$. According to the Euler equation, they spread this increase in their lifetime wealth into both consumption and bequest, based on their relative preference for them, defined by α.

In summary, the young generation in $t = 2$ is to inherit more than their parents did. Since the cost of education is fixed, in this period more are willing to become educated. The skill premium increases again, and the economy expands. Wealth and subsequent bequests of "unskilled dynasties" continues to rise, because of the dependence of current wealth on previous wealth. Similarly, both the "first generation" skilled agents and the "old rich" leave increasing bequests to their children. The modernisation of the labour force and of the economy continues. In a world with fixed population size, the process is bounded from above, reaching a steady state where everyone is skilled, and, over many generations, the distribution of bequest degenerates.

Along this development process, inequality is set to start increasing. In the beginning, the majority of the population remains poor, benefiting little from the modernisation process. Yet, the old and new rich gain considerably. Since new wealth is build upon old wealth, its expansion is faster for these groups, thereby increasing inequality. At some point however, this is reversed. More and more agents are becoming skilled, taking their families out of poverty. As these become the bulk of the distribution (i.e. the "middle class"), inequality starts to decrease. With fixed population, everyone is skilled, and wealth homogenises, eliminating inequality altogether.

1.5.2 A more formal treatment

The narrative of the development process can also be grasped from the model's equations. In particular, the development process hinges entirely on the behaviour of bequests. They give persistence, path dependence, and dynamic to the solution. As said, if they are increasing, the economy expands. If bequests converge, the model has a steady state; if bequests diverge or oscillate, there is no steady state.

The bequest that a young agent leaves when old depends on the group to which that young agent belongs (**NL**, **IB**, or **IL**), given by equations (1.18), (1.22), and (1.26) respectively. These can be represented jointly as a system:

$$x^o_{t+1} = x^y_{t+1} = \begin{cases} (1-\alpha)\left[(w^n + x^y_t)(1+r) + w^n\right], & \text{if } x^y_t < f_t \\ (1-\alpha)\left[(x^y_t - h)(1+i) + w^s_{t+1}\right], & \text{if } f_t \leq x^y_t < h \\ (1-\alpha)\left[(x^y_t - h)(1+r) + w^s_{t+1}\right], & \text{if } h \leq x^y_t \end{cases} \tag{1.39}$$

Each equation in (1.39) represents a first order, linear, non-homogeneous difference equation. While an algebraic solution for each of them could be computed, it is not the easiest method to analyse this system. Instead, a graphical study is more intuitive. Each equation can be plotted in the $\{x_t, x_{t+1}\}$ space, helping to understand the movement of bequests within dynasties, the educational status of agents, and the equilibriums at which bequests converge to or diverge from. Also, it is helpful to analyse both the short run and the long run of the model.

To set ideas, let us assume for the moment that the skilled wage is fixed (as in GZ93). Figure 1.8 (equal to Figure **??**, and Figure 1 in GZ93) plots the system in equation (1.39) (for clarity, the superscripts y are deleted from the bequest variables). Each solid segment in the figure corresponds to one equation of this system. The first one (**NL** agents) starts at a positive bequest level in $t+1$, has a positive slope of $(1-\alpha)(1+r) < 1$, and ends at a bequest of f in the horizontal axis. The second segment (**IB** agents) starts where the first segment ends, has a positive slope of $(1-\alpha)(1+i) > 1$, and finishes at $x_t = h$. The third segment (**IL** agents) commences at $x_t = h$, and has the same slope as the first segment. The 45 deg line represents those points where the bequest received in t and inherited in $t+1$ are equal.

Any point on the path represents the relationship between the bequest inherited by a young agent in t (the horizontal coordinate) and the bequest she leaves when old (the vertical coordinate), equivalent to the bequest inherited by her child, in $t+1$. Thus, this graph can be used to understand not only the change in the distribution of bequests in two subsequent periods, but to study it's evolution for the whole time span of the model. The direction of change is indicated by the arrows, on top of each line. Importantly, any point where a segment crosses the diagonal is an "equilibrium" point in the sense that a young agent inheriting such level of bequest would leave exactly the same amount to her child. In this particular example, there are three such equilibria. Following

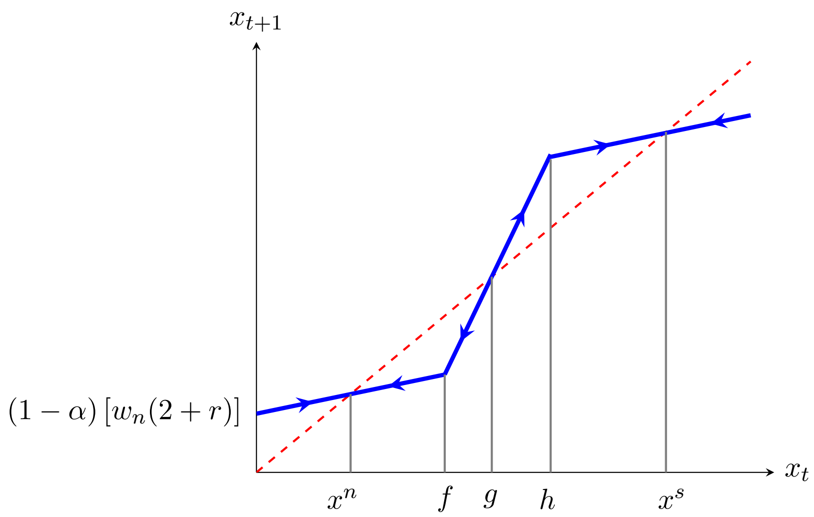

Figure 1.8: Dynamic of bequest between two periods, as in GZ93

GZ93 notation, these are x^n, g, and x^s respectively. Although not shown, the algebraic expression for each convergence point is found by imposing $x_{t+1} = x_t$ into equation (1.39) and solving for x_t.

Thus, according to the transition dynamics implied by the arrows, it holds that: dynasties where the young agent inherits a bequest smaller than g will move towards x^n; dynasties where the young agent inherits a bequest larger than g will move towards x^s; and those starting at g will remain there forever. In other words, knowing the initial distribution of bequest, we can predict what the final distribution of bequests is. In this particular case, bequests are falling for certain dynasties and increasing for other dynasties. Thus, this case is not enough to produce the Kuznets development process. First, the number of skilled is decreasing over time, because all those within f and g (**IB**) who become skilled leave lower bequests. In fact, in this economy, only those that start with $x^y > g$ will end up having skilled descendants in the steady state. This is true even if we bring the endogeneity of w^s back. Recall the level of the wage is given by the number of skilled. As this number is decreasing, the skilled wage is as well. Even more, as w^s falls, f_t (which is now time-varying) **increases** (see equation 2.8). This lowers L^s even further. In summary, the case in Figure 1.8, with or without endogenous wage, does not lead to the development process outlined earlier.

The scenario that leads to a Kuznets type of development process is shown in Figure 1.9. The central feature of this scenario is that **only one equilibrium for bequests is**

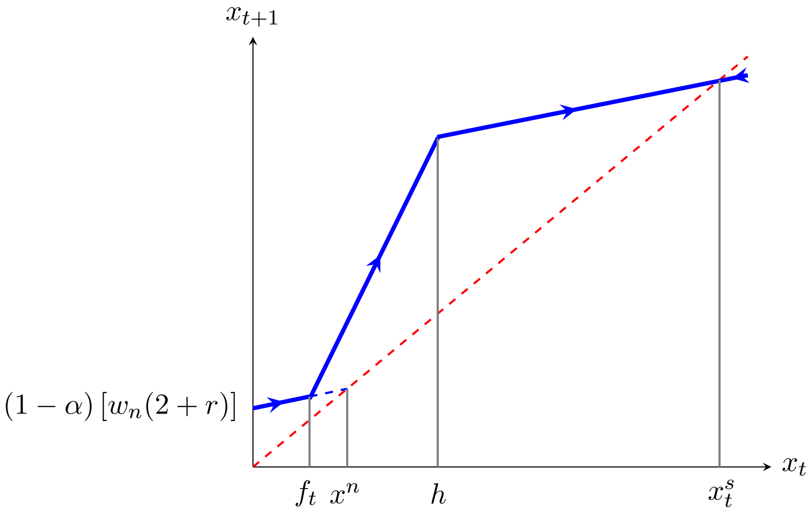

Figure 1.9: Dynamic of bequests for Kuznets development path

attainable, namely x^s_t (which is time-varying).[30] The other two equilibria are outside their corresponding intervals and thus are not attainable to agents. This is shown by the equations' projection outside the intervals, represented by the dotted blue lines (for example, $g_t < f_t$, inconsistent with the boundaries in equation (1.39)). Because only x^s_t is attainable, all dynasties move towards it. In particular, **all** those inheriting a bequest below that equilibrium leave increasing bequests, whereas those above that equilibrium move towards x^s_t by leaving lower bequests over time. Even more, as w^s is increasing over time (because more agents become skilled), f_t falls, further increasing the number of skilled agents every generation, and the subsequent bequest they leave. Over time, the economy moves toward higher L^s, until it reaches a steady state, where the economy is rich enough and all young agents choose education over work. Here, w^s and x^s_t stabilise at the level set out by $L^s = N$. After a certain number of generations, all bequests converge towards the steady state x^s, and inequality disappears.

Comparing diagrams in Figure 1.8 and 1.9 allows us to spot the central difference between them. Although x^s_t is attainable in both scenarios, x^n is only attainable in the former one.[31] As said, if both are attainable, bequests might increase for some

[30]More precisely, this is equal to $\frac{(1-\alpha)\left[w^s_{t+1}-h(1+r)\right]}{1-(1-\alpha)(1+r)}$.

[31]There are a set of **common** assumptions required for these scenarios to hold. First, both require **Condition 1**, which ensures that those with bequests higher than h are better of by investing in education. If this is not the case, there would be a continuous oscillation on L^s, from those who borrow to become educated, to their wealthy children who prefer not to invest in skills, whose children would then borrow to invest in education, and so on. More precisely, in Figure 1.8, bequests would oscillate between points g and x^s, with pivotal point h. In Figure 1.9, bequests would oscillate between points

dynasties and decrease for others – converging towards a bimodal distribution; if only x_t^s is attainable, bequests only increase – converging towards a unimodal, degenerate distribution. Mathematically, x^n is attainable *iff* $x^n \leq f_t$. Since f_t is a negative function of the skilled wage, the latter inequality holds when w_{t+1}^s is low enough. The intuition is the following: if this wage is low, the benefit of those who borrow to invest (**IB**) is not so great – yet still positive enough to make them become educated. Thus, their lifetime wealth does not increase by much, meaning **the bequest they leave to their offspring is not as great as the one they received**. This is clearly seen in Figure 1.8, as bequests of those along the $f_t - -g_t$ segment are below the 45 degree line. In turn, the following dynasties will also leave lower bequests, maybe not even investing in education. In other words, **education is useless as a social mobility tool**. The process continues until these dynasties converge towards x^n, and remain relatively poor forever. On the contrary, if the skilled wage is high enough, the **IB** groups is able to leave an inheritance higher than the one they received. In turn, their offspring also invest in education, continuing the upward cycle. This means **education becomes an effective social mobility tool** that lasts among generations. That is the central difference between these two scenarios.

Now, how low or high the skilled wage has to be depends of course on other parameters. In effect, it can be shown that $x^n \leq f_t$ is equivalent to:

$$w^n(2+r)\left[i-\alpha(1+i)\right] \leq \left[h(1+i)-w_{t+1}^s\right]\left[\alpha(1+r)-r\right]$$

Moreover, it can be shown that $[i-\alpha(1+i)]$, $\left[h(1+i)-w_{t+1}^s\right]$, and $[\alpha(1+r)-r]$ are positive (see footnote [31]). Thus, additionally to a low w_{t+1}^s, the aforementioned condition is more likely to hold for (i) high h, (ii) low w^n, and (iii) high α. This is, the higher the educational cost, the more agents in **IB** need to borrow, hence lowering the benefit of education and lifetime wealth. This makes children of first generation skilled less likely to pursue education. The same is true for both low unskilled wage and a low

f_t and x_t^s, with pivotal point h. Second, the stable equilibria – x^n and x_t^s – must exist. This is the case whenever the slope of the first and third path segments is lower than one, whereas the segment on the middle segment is larger than one. Yet, their existence do not mean they are attainable, as they could be outside the boundary of each path. Thus, the third and final common condition is that x_t^s is attainable. In practice, this means $h \leq x_t^s$, which can be shown to be equivalent to $h \leq w_{t+1}^s(1-alpha)$. Thus, as long as the skilled wage is high enough relative to the cost of education, bequests can growth past that cost. Of course, the smaller the proportion of lifetime wealth that is bequeathed by agents (given by $1-\alpha$) the higher the skilled wage has to be for this condition to hold.

rate of bequeathing (higher α). They lead to lower wealth accumulation, increasing borrowing costs of **IB** agents, and decreasing their lifetime wealth and that of their offspring. Finally, the net effect of r and i depend on other parameters.

Before moving on, let us bring the trivial equilibrium back onto the model. Recall from the previous section that in this economy there is **always** a trivial equilibrium, at zero skilled wage and employment. How does this fit into the above discussion? Well, if $w^s = 0$, x^s becomes unattainable, and only x^n is attainable. As such, bequests fall for those whose inheritance is above x^n. It is easy to show that in this case $x^n < f_t$ is always true. This means no one ever invests in education.[32]

1.5.3 Simulation

Having presented the particular scenario that gives rise to a Kuznets curve from both a narrative and graphical point of view, it is turn to simulate the model. The baseline calibration used is in Table 1.4. They meet the conditions outlined above, required for the scenario of interest to emerge (Figure 1.9. The weight that agents give to consumption in their utility functions (α) is assumed at 0.65. Thus, of every unit of wealth old agents have, they consume 65% of it, leaving the rest as a bequest. The annual interest rates are assumed 2% and 10% **per annum** for r and i respectively. Notice however that the model assumes each generation to last 25 years. Thus, the actual interest rates used in the model are their compounded values over 25 years. The level of variables like h and w^n have no particular interpretation. Their relative values are yet of importance. In this case, the cost of education is 27.5 times higher than the wage in the unskilled sector. Finally, the externality is assumed to be at 0.25, whereas the elasticity of capital is assumed to be 0.40%, in line with the empirical evidence of the labour share being around 60% of output (e.g. OECD, 2015). The size of the externality assumed yields increasing returns to scale of 1.25. This is relatively low, at least compared with other empirical studies evaluating human capital externalities (e.g. Arteaga Cabrales, 2010).

[32] The graph representing this situation is not shown, but it is constructed just as that of Figure 1.8. In particular, the middle path is short enough not to cross the 45 degree line at g_t, meaning the third segment does not cross that line either. This is, x^s is unattainable.

Table 1.4: Baseline calibration

Parameter	α	r	i	h	w^n	A^s	θ	ψ
Value	0.65	2%	10%	0.055	0.002	0.3	0.40	0.25

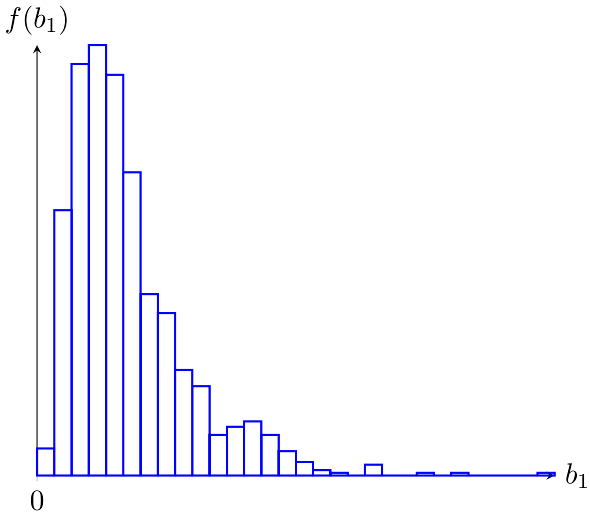

Figure 1.10: Initial distribution of bequests

Initial conditions

To start the dynamic processes of the model first it is necessary to populate the model with an old generation.[33] The economy starts at a "subsistence" level, where $L^s = 0$. The initial distribution of bequests **inherited by the young** in $t = 1$ (B_1) is drawn from a log-normal distribution with parameters $\mu = 0.03$ and $\sigma = 0.65$, which leads to an initial wealth inequality of 0.36. This distribution is shown in Figure 1.10. This also includes the cost of education h, showing that those who are rich enough to invest without borrowing are 11%. Naturally, since all agents are unskilled workers earning w^n, wage inequality is zero.

Results

The model is simulated for 7 periods (enough for convergence). Simulation results are presented in Figure 1.11. The first three graphs show the production inputs, capital,

[33] In practice, the simulation is carried out for a **generation** size of 1,000 (population size of 2,000). This is however irrelevant, as everything can be scaled up or down (e.g. assuming one dynasty represents m dynasties). In turn, the employment variables are presented as proportions with respect to the generation size.

unskilled labour and skilled labour. Recall that the size of every generation is 1,000, and that both young and old can work in the unskilled sector. There is convergence of these inputs to their steady state level between periods 4 and 7, which given our 25 year generation definition, imply a transition between no development to complete development of at least 100 years. Of particular interest is the skilled market. The economy started without skilled workers. After the first generation, around 28% of the young agents decide to become skilled. This reaches 76% in the next generation, until everyone become skilled by the 6th generation. As expected, the skilled wage follows a similar pattern, until it stabilises at 0.56. Once the economy is fully developed, the skill premium with respect to old technology is 280.[34] Per capita GDP grows at 16% *per year* during the first generation (when the economy moves from a traditional economy to a developing one). The next 25 years the annual growth rate is of 5.8%. Once the economy has reached the 88% of its long run GDP per capital level, the rate of growth slows down to 1.4% per year, until it becomes negligible in the last generations. Finally, Gini for wealth and income increases in early stages of development, while falling in later periods. The change in income inequality is much more sharp. This is because, unlike bequests, wages have no path dependence.

Having simulated the model, Figure 1.12 presents the resulting Kuznets curve, based on both wealth and wage. The former is calculated over inheritance, and thus it is at the level of the dynasty or household, whereas the latter is at the individual level, including both working young and old agents, even if from the same dynasty. The lack of path dependence in wages also imprints a dramatic evolution in the Kuznets curve under this definition. Finally, in this particular setting without long term growth or population growth, the Kuznets curve becomes vertical, a clear artefact of the fixed population and exogenous unskilled wage assumptions.

Middle class

Let us reconnect the analysis with the study of the middle class. Figure 1.13 presents six graphs, with different perspectives on the middle class, related to the discussion in

[34]This is surely the most unrealistic result of the simulation. Yet, this feature of the model is largely an outcome of the static unskilled sector where the wage is exogenous. It would be possible to endogenise w^n by adding spillovers between sectors, or demand factors that make agents want to consume goods from the traditional sector. See later for a brief discussion on this.

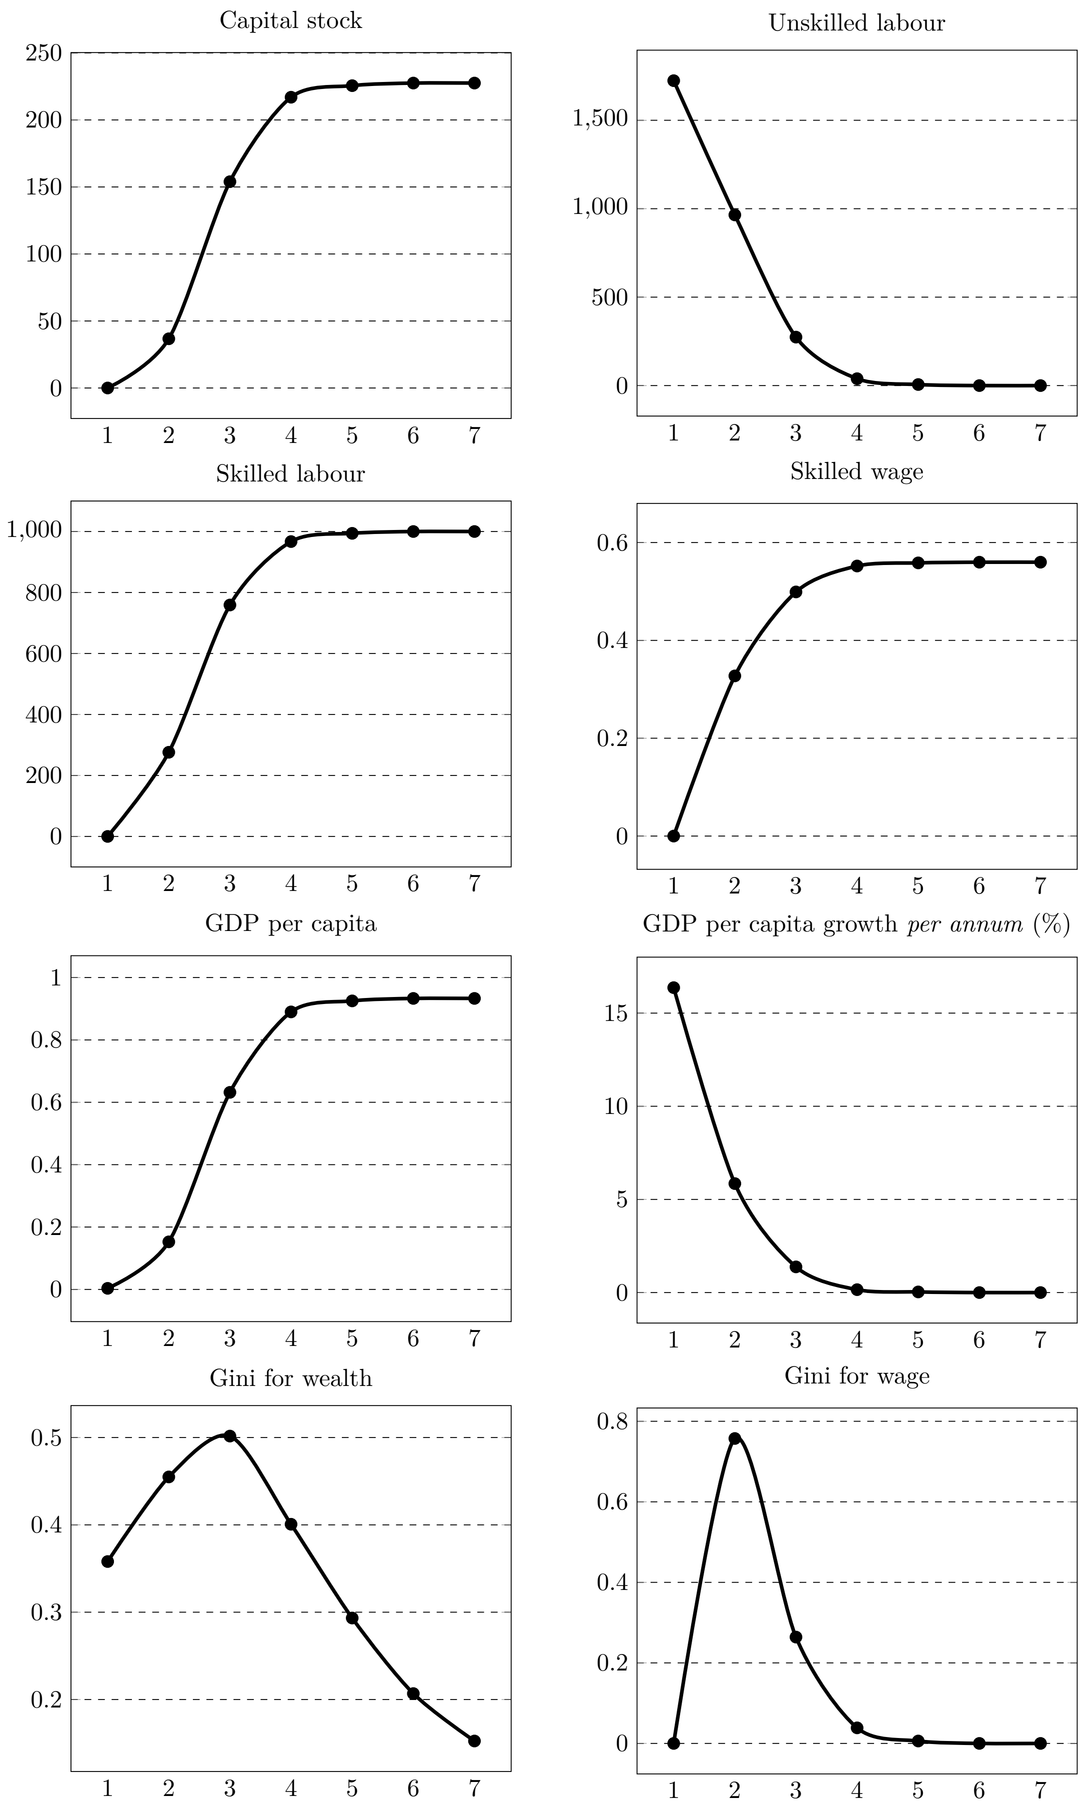

Figure 1.11: Evolution of key variables in the model, baseline calibration

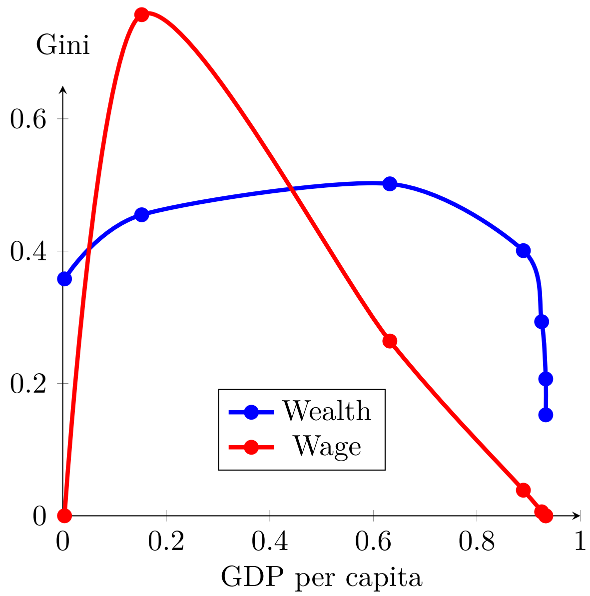

Figure 1.12: The Kuznets curve for wealth and wage, baseline calibration

the introductory chapter. Notice that the convergence of bequests takes longer than the convergence of production factors. The latter stops as soon as everyone become skilled. In consequence, more periods are added to the analysis.

The top two graphs analyse the increase in average wealth of the middle portion of dynasties, ordered by wealth level. These correspond to the definition of middle class as a fixed proportion of the population, and the thresholds are equivalent to those in Atkinson and Brandolini (2011) and Easterly (2001) respectively. It is clear that the wealth accumulated by the middle class in the simulation grew considerable over the period. A fairly similar result is obtained using the definition of the middle class as a variable portion of the population, shown in the middle panels. The thresholds used are related to the median wealth in the population, as in Birdsall, Graham and Pettinato (2000). These results show an initial increase after the first period, but a fall in the next one, probably due to the large increase in skilled labour between periods 2 and 3 (from 276 to 759), pushing median lifetime wealth up. Yet, this trend is quickly reversed, and a greater proportion of the population is classified as middle class. Finally, the bottom panels focuses particularly on the top, showing that the share of wealth owned by the top dynasties decreases consistently over the period (except for a spike in the first period, due to some agents becoming skilled). All the above highlights the fact that, as long as the economy develops, and a Kuznets curve arises, the middle class is "rising", regardless of how it is defined.

Departures from baseline calibration

As it turns out, the alternative dynamics of the model are pretty straightforward. As long as x^s_{t+1} is the only attainable equilibrium in the economy, there is full transition from the traditional towards the modern sector ($L^s = N$). The only change is the rate of transition towards it. For example, an increase in ψ – the externality – leads to a quicker transition path. Similarly for a lower educational cost. The effect of the unskilled wage is interesting. On the one hand, a higher w^n lowers the incentive to switch. Yet, it increases the accumulation of wealth, leading to larger inheritances. The result is a "twist" in the skilled labour path such that less agents invest in the first periods, but the steady state is reached quicker. On the contrary, a lower w^n produces a higher initial increase in L^s, whilst delaying the development process. This is because the accumulation of wealth requires more time in order to reach levels good enough to motivate these dynasties to borrow and become skilled.

The role of θ is also intuitive. The higher this is, the slower the development process. This is because the elasticity of labour $(1 - \theta)$ directly affects the marginal product, and therefore the skilled wage. The role of α is also evident. The higher the proportion of lifetime wealth old agents leave to their offspring, the faster the transition towards complete development. Additionally, a shrink in the interest rate spread also increases the speed of transition, as the cost of borrowing relative to the benefit of savings falls.

Finally, the role of the initial distribution of bequests is relevant, but not as much. In effect, even if all agents start with zero wealth, the model converges. Again, this is an outcome of x^s_{t+1} being the only attainable equilibrium. Naturally, the development process might take longer, because wealth accumulation starts from zero.[35]

Now, any change in the parameters that introduces x^n as a possible equilibrium alters the results. Here, the economy moves towards $L^s > 0$, either reaching a steady state

[35] In this extreme example, all agents are always equal, meaning full development is achieved in one generation. Inequality is always zero.

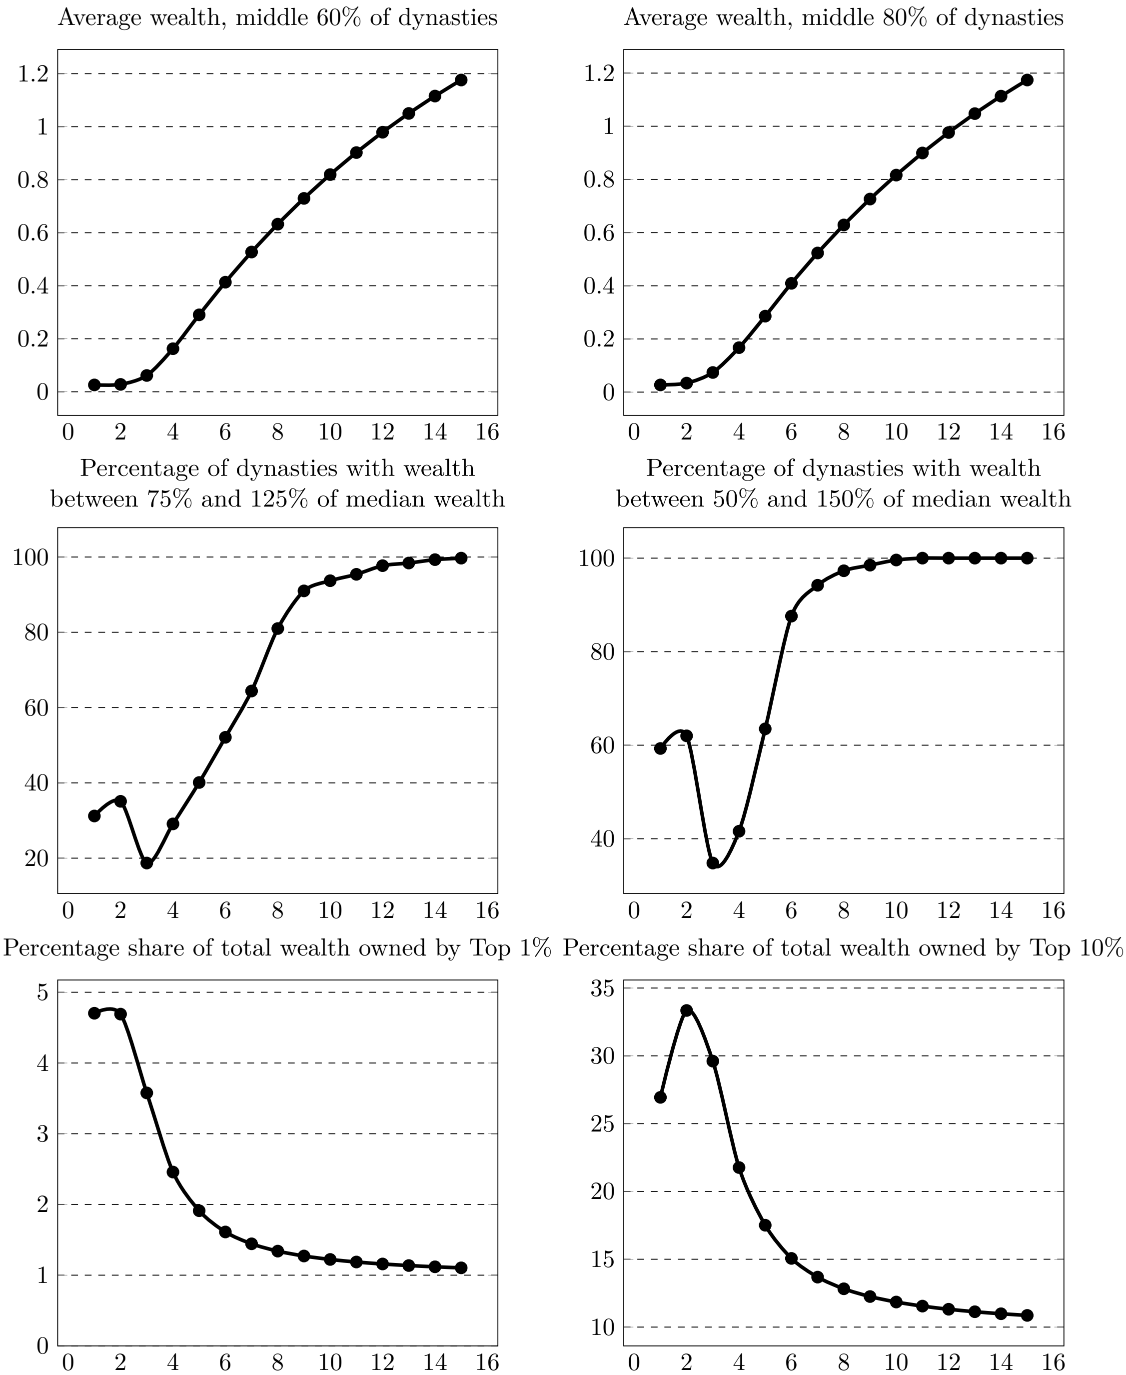

Figure 1.13: Evolution of the middle class among dynasties using wealth (bequests), baseline calibration

with positive skilled labour, or, if the conditions are not favourable to sustain the latter, a reversion to the full traditional economy. This is the case for example if the externality is very low. The economy ends up in the trivial skilled labour market equilibrium.

In summary, the parametrisation defines the steady state the economy achieves. The Kuznets curve is a rather natural outcome of an economy where only x^s_{t+1} is attainable, and many parameter combination would lead towards it. Some exceptions are, as mentioned, (i) when all agents are closely equivalent in terms of initial wealth, or (ii) a very quick convergence (which does not capture transition processes).[36] Alternatively, if x^n is also attainable, the economy might end up at an intermediate level of development (as in the standard GZ93 setting), or remain undeveloped forever.

Causality

The causality between inequality and development implicit in the model is not as clear as it seems. In the one hand, for the same mean of the wealth distribution, a higher variance leads to slower growth, as less agents invest in education. Notice that the story about the rich saving more than the poor, which then boosts investment plays no role here because the economy is open to the capital markets. Thus, there is no **direct** effect of inequality on the capital stock, but only through L^s (recall equation 1.8). The effect of development on inequality is obscured by the bequests dynamics. On the one hand, growth is itself a product of the equilibrium in the labour markets. Yet, the higher this equilibrium is – and thus the rate of growth, the faster bequests grow. Now, since the rate of growth is higher in the first periods, it means the dis-equalisation process is faster when the economy grows faster. When the economy is growing little, its because bequests are converging, reducing inequality. This suggests a negative correlation between inequality and GDP growth. Still, causality is not clear because the central driver of the model is wealth accumulation, which drives both inequality and growth.

[36]The former is certainly possible in real economies, whereas the latter is probably not very realistic.

1.5.4 Exogenous growth

The model used here employed a form of endogenous growth which leads to a "total" aggregate demand for skilled labour to be upward sloping, albeit fixed over time. An alternative approach would be to introduce exogenous growth in the production function. For example, instead of using equation (1.4), technology could be defined as $A_t = A^s(1+g)^t$, with $g > 0$. In this case, the total demand for skilled labour would be perfectly elastic (which can be seen by replacing $\psi = 0$ in equation 1.10). Yet, it would move up every time because of the exogenous increase in labour productivity. This is equivalent to our constant but upward sloping demand curve with endogenous growth. In effect, for a certain calibration it is possible to find a path of skilled labour equilibria equivalent to that in Panel (a) of Figure 1.12. There are however very important differences. First, the driver of development is by definition exogenous, whereas in the model here it is explicit. In fact, the previous section highlighted the crucial role the level of the externality plays on the possibilities the economy has to become rich. Any such insights are unavailable from a generic growth mechanism. Second, exogenous growth always leads to full development, regardless of any parametrisation (given that **Condition 1** holds). This is because the skill premium permanently widens every period – even after $L^s = N$. This determinacy in the development outcome already neglects any interesting insight that arises from alternative parametrisation, providing no understanding of the conditions under which partial or no development at all might occur. Third, the immediate cause for economic transition in the model with exogenous growth is the skill premium. In our setting however, the immediate source of growth is wealth accumulation. Namely, it is the continuous increase in dynastic wealth what drives more young agents to become educated. The increase in the skill premium – endogenous – is merely a reflection of this. Actually, in our setting, the skill premium reaches an equilibrium when $L^s = N$. In summary, endogenous growth adds flexibility to the model's outcomes whilst enabling much richer insights from its mechanisms. This is, in our view, enough to justify the choice.

1.6 Conclusion

This chapter presented, solved, and simulated a model where an externality from the acquisition of skills by agents results in increasing returns to scale in the skilled sector of the economy (i.e. SBTC), leading to a positive spiral of growth and a full modernisation of the economy. Along that process, inequality follows a Kuznets shape. Still, the model itself is flexible enough to produce anything between zero to full development of the economy.

The role of the externality is of particular importance. An economy might never develop fully if the externality from knowledge is not high enough. This highlights the importance of spillovers from human capital into the wider economy. It is of the interest of governments to spur the conditions that maximise the externalities from knowledge, as well as to facilitate the access to human capital by agents, in particular if these are financially constrained, or the uncertainty about its benefits is high. This model highlights that underinvestment in education can well be a deep factor behind underdevelopment.

Another central result of the analysis is the crucial role of social mobility coming from education. It was shown that only when education becomes a social mobility tool within dynasties (i.e. that children of educated parents also become educated) the economy can flourish over time. If that is not the case, there is never complete transition from underdevelopment towards development.

The use of an endogenous growth model also highlights the intrinsic endogeneity of the skill premium. Instead of simply assuming an ever increasing skill premium, this is a natural outcome of the development process. This shifts the causality of development in important ways. In effect, it is not the skill premium which drives development but the accumulation of wealth that enables agents to become educated.

Naturally, there are several short-comings of the model, which could be relieved by adding further complexity to it. Two extensions that have been partially explored but are not included here are poverty traps and taxation. The former arises by limiting the capacity to save for those with a low enough wealth, either because of subsistence levels of income or credit constraints. These households cannot accumulate wealth and might never transition out of poverty. Taxation can be introduced via an inheritance tax, a lump sum tax, or VAT. It can be shown that the first two are sub-optimal because they slow down the wealth accumulation process (unless the receipts are used to provide public education). On the contrary, a VAT does not alter the incentives to leave bequests.

Other extensions were already mentioned in Section 1.2. The distribution of power via wealth could prove to be interesting, which might be linked to taxation issues mentioned above. Similarly, endogenous fertility might introduce interesting constraints into the development process by directly affecting the process of wealth accumulation. For example, the more children old agents have, the lower the bequest each of them get. This might be a particular problem for poorer dynasties, maybe restricting their upward social mobility permanently.

It is also of interest to explore information problems, and uncertainty. For example, in a more general setting where workers are not fully rational, information is imperfect, or there are uncertainties, it is likely that many workers underinvest in their education. As said, a government interested in development would aim to overcome this, and exploit the human capital's externality in full, either by providing (compulsory) public education (leading to the skilling of the workforce), or loans, both financed through taxation.

Probably the most profitable extension however would be to make the unskilled wage endogenous. One mechanism would be through positive spillovers from the modern to the traditional sector. This would make the convergence toward full development much slower, probably requiring a larger externality in order to achieve full modernisation. Another mechanism is through demand channels. For example, instead of being two technologies for one good, they could be two goods. As Mazzolari and Ragusa (2013) show, a large proportion of the increase in low-skill jobs in the US can be explained from the increase in demand for time-intensive services – mainly produced by low skill workers, by an expanding skilled workforce. Thus, by assuming Y^n to be services and

Y^s to be goods, together with imperfect substitution between them in agents utility it is possible to partly off-set the increase in the skill premium and the consequent complete specialisation of the economy. It is expected however that a Kuznets curve would still arise in this setting. Notice that in this chapter inequality falls in later development stages because the majority of the population are becoming skilled. In some economies however, inequality might fall because the skill premium is partly eroded, without need of an increasingly homogeneous labour force. As such, turning the unskilled wage endogenous might not necessarily undermine the analysis in this chapter, but would actually enrich it by introducing another equalisation source.

All in all, the model presented here contributes to the literature by providing a micro-founded Kuznets curve that arises under a reasonable development process, whilst providing interesting insights about the factors that lead toward full, partial or no development at all. In the near future I expect to dig deeper into these factors by expanding the model towards further realism and relevance.

2

Job Polarization and the Distribution of Skills

2.1 Introduction

It is well documented that advanced economies have experienced shifts in employment from mid-skill intensive jobs to **both** low and high skill intensive jobs, at least since the 1980s (Autor, Katz and Kearney, 2006; Acemoglu and Autor, 2011; Goos, Manning and Salomons, 2014). This hollowing out of the labour markets is known in the literature as job polarisation. Figure 2.1 presents further evidence of this, for some of the largest EU countries, using harmonised employment data from Eurostat. In particular it shows the change in employment shares for three occupational groups, between 1992 and 2008. The case of polarisation is clear in several countries like Denmark, Germany, the Netherlands, and particularly the UK. In other countries there is also a considerable fall in mid-skill employment, but employment reallocation is mainly unidirectional, for instance, up in Italy and Spain, and down in France.

The above evidence give rise to two, related questions. First, what has driven job polarisation? Second, why some countries experience alternative forms of employment reallocation? The theoretical literature has proposed several drivers of job polarisation in general, and of reduction in middle skill jobs in particular. Perhaps unsurprisingly, technological change emerges as prime candidate. Moving away from the standard

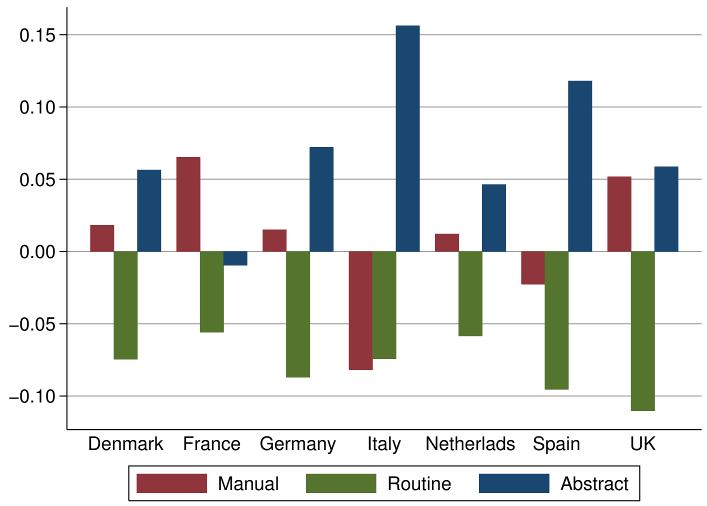

Figure 2.1: Change in employment shares, 1992-2008, selected EU countries

but inflexible skill-biased technical change (SBTC) hypothesis – which cannot alone explain the observed employment polarisation trends because it focuses on two workers categories (high and low skill), several authors have explored a more refined form of technical change, called task-biased technological change (TBTC). This literature is based on Autor, Levy and Murnane (2003), who shift the focus on "tasks" (type of job activity) rather than skills (type of workers). By differentiating between routine, manual and cognitive tasks, TBTC is based on the idea that new technologies are complementary with routine tasks (which evidence indicates correspond to middle-skill jobs), reducing the demand for workers on these jobs. For example, Acemoglu and Autor (2011) present one model where capital is solely a producer of routine tasks, hence allowing firms to replace workers with capital in that occupation. Autor and Dorn (2013) follow a similar approach, where capital only contributes toward routine tasks production. To add dynamics into the model, they assume the price of (computer) capital falls exogenously at an exponential rate. In both models, increasing efficiency or cost of capital induces employment reallocation away from routine jobs. This effect of technology at the task dimension is what defines TBTC.

Another channel by which technological change can lead to polarization is off-shoring. Since technical improvements enhances the reliability and decreased the cost of ICT, trade costs fall. In response, firms shift sections of their production into countries with cheaper labour, thereby producing *trade in tasks* (e.g. Grossman and Rossi-Hansberg, 2008, 2012). Under simple parametrisations, it is possible for routine tasks to be off-shored, inducing greater comparative advantage of the advanced countries in both manual and cognitive tasks, which yields job polarisation.

Job polarisation is also thought to be driven by changes in the demand for goods. For example, Mazzolari and Ragusa (2013) show that the increase in skilled workers has increased the demand for time-intensive services – mainly produced by low skill workers. They show that this channel explains a considerable part of the observed increase in low-skill jobs in the US. Alternatively, Moreno-Galbis and Sopraseuth (2014) propose the change in population ageing patterns as a force behind the larger demand for low-skill services. In particular, they modify Autor and Dorn (2013)'s model by adding two types of workers, young and old, with different preferences, where they show how an increase in the proportion on old agents (i.e ageing) raises the demand for low-skill services.

Several empirical studies put these hypothesis to the test, finding different degrees of support for them, depending on countries and periods (Goos, Manning and Salomons, 2014, Oldenski, 2014, Nellas and Olivieri, 2012, Firpo, Fortin and Lemieux, 2011). However, what it is clear from this literature is that the focus has been mainly on the demand side of the labour market, with complete disregard to the supply of skills. For example, can differences in the level and changes in the ability of workers help explain the different patters in Figure 2.1? More precisely, what is the role of heterogeneity of skills in both the degree of middle-skill job destruction and the direction of reallocation?

The model in this chapter explores these issues (Chapter 4 takes the model to the data). The core of the model is workers' ability for different tasks (manual, routine, and abstract), which are assumed to be heterogeneous. In particular, ability in the population distributes log-normal, assumption justified with the empirical analysis of Chapter 3.

Since workers are fully mobile across tasks, they sort into tasks (occupations) based on comparative advantage. Furthermore, the model assumes absolute advantage of ability among workers. This is, if a worker is more able than another at a given task, she is also more able at all other tasks. It is shown that this assumption allows for a tractable sorting equilibrium, where workers with the highest ability are employed in abstract tasks, and those with the lowest ability are employed in manual tasks.

The production framework of the model has two layers. First, workers sort into tasks, defining the employment share on each occupation. Then, workers on each task contribute to the production of a task output, equivalent to the sum of employed workers' ability. Thus, the more able workers allocated to a given task are, the higher the output from this task. Then, these outputs are used by firms to produce the final good Y, using a CES technology (there is no capital). Firms operate in a closed economy with competitive markets.

Besides this being a very simple framework, the model yields no closed form solution for employment shares. However, enough intuition arises from a graphical analysis, which is used to understand job polarisation. Additionally, it is still possible to provide a comprehensive analysis of the role each parameter has on the dynamics of employment, via the Implicit Function Theorem. This approach is used to characterise two distinct, exogenous sources of job polarisation: (i) changes in the distributions of ability (the novelty with respect to the literature), and (ii) task-based technological change (TBTC). The latter can be easily introduced into the model as the effect of changes in the idiosyncratic productivity level of each task on the production of the final good Y. This flexible specification allows for any type of TBTC to be analysed; for instance routine biased, manual biased, etc.

Several insights arise from these exercises. First, the elasticity of substitution between tasks (σ_Y) plays a pivotal role in most of the exogenous changes that might affect the economy. For instance, if this substitution is relatively high ($\sigma_Y > 1$), firms will expand the use of any input that has become more productive (either because of workers ability or technical change). Conversely, low substitution ($\sigma_Y < 1$) induces firms to release resources from more productive tasks into those who have not benefited from changes in ability or technology. In the case of a Cobb-Douglas production function ($\sigma_Y = 1$), no reallocation is required.

Second, polarisation cannot arise solely by changes in the ability of workers for manual **or** abstract tasks alone. For instance, a general increase in the ability of workers for abstract occupations will induce employment reallocation "down" to routine and manual tasks (if $\sigma_Y > 1$) or "up" from manual and routine tasks (if $\sigma_Y < 1$). Recall that job polarisation requires movements "down" (routine to manual) and "up" (routine to abstract). This unidirectional reallocation is of course a consequence of the simple sorting mechanism, based on absolute advantage. Polarisation might arise however from a sole change in workers' ability for routine tasks. Additionally, it is shown that simple changes in comparative advantage alone cannot induce job polarisation. Naturally, many combination of parameter changes lead to the same result, albeit not in generalisable ways.[1]

The results from the analysis of technical progress mimic those of ability. Namely, an increase in productivity of manual **or** abstract inputs in production cannot lead to polarisation, as they only generate reallocation "down" or "up". Yet, changes in routine input's productivity alone can produce polarisation. Combination of these factors also might induce polarisation.

Interestingly, these results are very much in line with the original Autor, Levy and Murnane (2003) argument that TBTC induces automation of routine tasks, enhancing their productivity. Regardless of whether computers increase the productivity of workers at routine activities or enhances the productivity of routine inputs in production (or both), this is consistent with our argument, provided the elasticity of substitution among tasks is lower than one. This is, as (workers at) routine tasks become more productive, firms demand less of this input, inducing reallocation of employment away from L_r and into L_m and L_a, i.e. job polarisation. As Chapter 4 discusses, the evidence seems to be very much in favour of low elasticity of substitution.

This chapter provides several contributions to the literature. First, it presents a model that is novel in the literature. Its addition with respect to the current sorting literature is the explicit use of log-normal distributions. The analysis of absolute and comparative advantage arising from such distributions are relatively straightforward, and the model allows for considerable intuition regarding sorting. Additionally, the framework is very easy to generalise, as Appendix F shows.

[1] This is, the polarising outcome can only be understood on a case by case basis, depending on the particular parametrisation used.

Second, the model adds a channel leading to polarisation that so far has received little or inadequate attention in the literature, namely the role of workers' ability. Most of the existing literature (theoretical and empirical) tends to focus on technical change (like ICT, TBTC, offshoring, etc), the role of demand, or on institutions (see Section 2.2 for more details on the literature). By contrast, it is shown here that changes in workers' ability alone can lead to polarisation, without need for TBTC. More generally, changes in workers' ability can lead to any reallocation pattern, and certainly to all those presented in Figure 0.4, presented in the introductory chapter. Thus, the framework used here brings workers' ability, and more generally the sorting of workers and comparative advantage, to the core of the analysis. The complementary empirical analysis in the next two chapters further supports the empirical relevance of the model.

Finally, the model is built upon relatively robust foundations, in contrast with some existing literature. For example, Autor and Dorn (2013) develop a model of polarisation where workers have homogeneous ability for manual and abstract, and their ability for routine tasks is exponentially distributed. Additionally, they do not allow for mobility between routine and abstract tasks. As Chapter 3 shows, all these assumptions are very much at odds with the evidence. Conversely, the model in this chapter builds upon stronger empirical foundations (again, see aforementioned chapter), whilst still providing a clear framework with which to study job polarisation. Additionally, the model is readily applicable to empirical analysis, as the rest of this thesis shows.

Road map

This chapter continues as follows. Section 2.2 provides a survey of the relevant literature. Section 2.3 presents the complete model. Section 2.4 describes the general solution of the model, and studies its key markets. Section 2.5 applies the model to the study of job polarisation. The final section concludes.

2.2 Literature review

This chapter is connected to two areas of the literature. On the one hand, its aim relates to models explaining recent trends in job polarisation. On the other hand, its methodology relates to models assuming relatively complex patterns of skill heterogeneity among workers, in the context of *tasks*. In this section I review both lines of the literature.

Job polarisation

Many authors have documented the apparent polarisation of employment and wages in advanced economies. Acemoglu and Autor (2011) document the fall in the *share* of employment of middle-skill occupations – which they define as “comprising sales; office and administrative support; production, craft and repair; and operator, fabricator and laborer” (page 1045) – and the corresponding increase in low and high skill employment shares for the US in a window of 30 years, results that are robust to different occupational aggregations and data sources. Similarly, Goos, Manning and Salomons (2014) show the polarisation trend in employment shares in 16 European countries for the 1993-2006 period, with a detailed analysis per industry and different occupational categories.

Technological change arises as a prime candidate when trying to explain these patterns. However, the standard view of technological change up to the last decade – that of skill-biased technical change, or SBTC – cannot alone explain the observed employment polarisation trends because it focuses on two workers categories (high and low skill). Indeed, SBTC argues that technology is mainly complementary to high skill workers, therefore raising its demand and relative wage, and leading to higher inequality. This is not consistent with the observed increase in low skill employment shares.

In order to properly understand polarisation, a more refined view of technological change was needed. This new perspective is called task-biased technological change (TBTC). It centred in the concept of tasks, usually defined as “a unit of work activity that produces output” (Acemoglu and Autor, 2011). The methodological “innovation” is the use of a multilevel production function, where workers capabilities or attributes contribute to an intermediate input – here, tasks – which then are combined to produce the final good.[2] Then, unlike SBTC that affects worker types, TBTC affects relative

[2]This multilevel approach to production is of course not new. Early contributions include Hartog (1979) and Sattinger (1980). Similarly, the use of tasks as intermediate inputs in production pre-dates the recent literature. A prime example (reviewed later) is Sattinger (1975). The availability of the Dictionary of Occupational Titles in the US since 1938 has also allowed tasks to be made operational, allowing empirical research. Again, Hartog (1979) and Sattinger (1980) are examples of this.

productivities between tasks. By assuming a particular mapping between workers and tasks it is possible to match the facts. For example, Autor, Levy and Murnane (2003) categorise tasks as routine and non-routine, arguing the former are substitutes with computers whereas the latter are complementaries with computers. Then, they show how the observed increase in the use of computer capital by firms (proxied by hardware and software spending) explains part of the fall in the routine content of jobs. While Autor, Levy and Murnane (2003) do not study job polarisation, nor advance the TBTC hypothesis, the connection is evident. By noting that middle-skill jobs tend to be routine-intensive whereas low-skill and high-skill jobs are mainly based on non-routine tasks, technological change – via computer capital and automation – can lead to a reduction in routine jobs and a corresponding increase in other job categories, either due to complementarity or simple resources' reallocation. This effect of technology at the task dimension is what defines TBTC.

There are many models where technological change is the leading cause of job polarisation. Acemoglu and Autor (2011) present one model where capital is solely a producer of routine tasks, hence allowing firms to replace workers with capital in that occupation. Autor and Dorn (2013) follow a similar approach, where capital only contributes toward routine tasks production. To add dynamics into the model, they assume the price of (computer) capital falls exogenously at an exponential rate. The polarisation result arises whenever the elasticity of substitution in production is low compared with the elasticity of consumption (more on this model later).

Another widely explored channel of how technology can lead to polarization is offshoring, which refers to the reallocation of a certain production process by firms to foreign countries. Given that technological change also improves the reliability and decreased the cost of ICT, trade costs are in effect lowered. In response, firms find it profitable to shift sections of their production (manufacturing, services, R&D, etc) into countries with cheaper labour, thereby producing *trade in tasks* (e.g. Grossman and Rossi-Hansberg, 2008, 2012).

Certainly, technology is not the only cause of polarisation proposed in the literature. Changes in the demand for goods seem to be an attractive candidate too. For example, Mazzolari and Ragusa (2013) state that the increase in skilled workers has increased the demand for time-intensive services – mainly produced by low skill workers – because these, otherwise home-produced services, are now bought in the market (e.g. substitution of home meal production for take-aways or eating out). They show that this channel explains a considerable part of the observed increase in low-skill jobs in the US. Alternatively, Moreno-Galbis and Sopraseuth (2014) propose the change in population ageing patterns as a force behind the larger demand for low-skill services. In particular, they modify Autor and Dorn (2013)'s model by adding two types of workers, young and old, with different preferences, where they show how an increase in the proportion on old agents (i.e ageing) raises the demand for low-skill services.

Evidence is rather mixed regarding the different hypotheses, but in general there is evidence in favour for all of them (Goos, Manning and Salomons, 2014, Oldenski, 2014, Nellas and Olivieri, 2012, Firpo, Fortin and Lemieux, 2011). However, what it is clear from the literature is that so far, heterogeneity of skills among countries has not been regarded as important to understand job polarisation, very likely because there has not been a purpose-built model to highlight the role of skills. That is the goal of this chapter.

Heterogeneous labour

The model I develop in this chapter is one with heterogeneity of skills among workers. Since this is a core element of the model (for which evidence is quite supportive, according to results in Chapter 3), it is interesting to see how it relates to the rest of the literature. The heterogeneous labour literature is of course huge, and a survey of it is well beyond the scope of this chapter. My particular focus here is for those models that use *tasks* as the foundation of their labour framework. As stated before, the concept of tasks as a modelling layer in production is rather old in the literature. Yet, until recently, it was not at the centre of the technological change literature. Autor, Levy and Murnane (2003) reversed this, and the number of studies using tasks has proven quite large since then.

The key concept of tasks is its definition as an intermediate input, which itself is produced by primary factors like capital and labour. This means an otherwise standard production framework is now multilevel. Since it is usually assumed that workers can perform many tasks (otherwise, there is little benefit from this approach), a sorting element arises. Hence, in many of these models, the standard problem to solve is about allocating workers into tasks.

The simplest case is that of a model with homogeneous workers. For example, Grossman and Rossi-Hansberg (2008) consider two types of workers, L and H. Whereas L differ from H, there is no difference between two L-type workers (or two H-type workers). This means that whatever changes for L (wage, utility, etc), changes for all workers in L. The implications of a framework like this are strong (e.g. all workers benefit, all have incentives to reallocate, etc).

Alternatively, heterogeneity enriches the results. For instance, Acemoglu and Autor (2011) assume three types of workers (H, M, L), which allocate into a continuum of tasks according to comparative advantage. Workers differ in their ability for each task. More precisely, defining the continuum of tasks over the i dimension, $i \in [0, 1]$, ability of H-type workers is $\alpha_H(i)$, which varies over i. Same for the other workers. The particular assumption about $\alpha_j(i)$ determines workers' comparative advantage and their allocation into tasks. For our interests, the main problem of this framework is that what defines the routine or non-routine content of tasks is not a characteristic of tasks themselves but of the worker types. For instance, Acemoglu and Autor (2011) identify as routine tasks those performed by M-type workers. Although simple, this means that the supply of workers is fixed in advance, i.e the model does not generate reallocation of workers among tasks but reallocation of tasks among workers. Therefore, this framework is less useful in explaining job polarisation patterns because worker's classes are fixed *ex-ante*.

Other models assume workers heterogeneity by starting from a unique distribution of skill among workers, which then is mapped into particular abilities for several tasks. For example, Yeaple (2005) assumes an economy with three occupations (Y, H, L) and a continuum of workers, whose ability distribution is $G(Z)$. The mapping from ability Z to productivity is given by the functions $\varphi_j(Z)$, for each occupation j. Clearly, the assumptions behind the latter functions are essential to determine who works where. Yet, the particular assumptions they make have no simple interpretation in terms of common distribution, nor do they focus on understanding the consequences of their assumptions. In my model however – besides having a completely different production framework –, assumptions are carefully chosen with respect to the underlying distributions that result, in light of the empirical evidence this thesis produces, presented in Chapter 3.

Autor and Dorn (2013) is another interesting example, and partly the inspiration of the model used here.[3] They model an economy with two types of workers, low skill (L) and high skill (H), and three occupations or tasks: "abstract", "routine", and "manual". The central difference between L and H is that they perform alternative tasks. On the one hand, L-type workers can perform either the routine task or the manual task. On the other hand, H-type workers can only undertake the abstract task. Thus, from the labour market point of view, the interest is only in how low-skill workers sort between manual and routine tasks. They assume that the ability level of L-type workers in performing the manual task is homogeneous whereas their ability in performing the routine task is heterogeneous. Intuitively, the most skilled routine workers allocate to that task, whereas the rest sort into manual ones. The model in Autor and Dorn (2013) is yet more complex. They have two sectors – goods and services, where capital can also produce routine tasks, which are then used as inputs for goods production. Conversely, manual tasks can only produce services. Thus, in essence, low-skill workers mobility across tasks also implies mobility across sectors. Assuming the price of capital falls exogenously, they show that polarisation arises whenever routine labour is relatively easy to substitute for capital, and consumers do not easily substitute goods and services. This is, mobility of workers becomes an exercise of fine-tuning the best combination of final production and intermediate production in terms of utility.

Clearly, Autor and Dorn (2013) model is complex and asymmetric both in its production and labour sides, and limited in terms of mobility and heterogeneity. The model developed here can be thought of as an extension of theirs. First, it introduces full mobility of workers among tasks. As Section 3.6 in Chapter 3 shows, mobility between routine and abstract tasks is empirically important. Second, it allows for heterogeneity in workers ability in all tasks, instead of just one. The cost of this is added analytical

[3]There are many other interesting papers in terms of labour framework within tasks but many of them are too complex for the interest here (like Costinot and Vogel, 2010, who works with both a continuum of workers and a continuum of tasks, or Grossman, Helpman and Kircher, 2013, who add heterogeneity within a 2x2x2 trade model).

and computational complexity. In return, the production framework in this model is much simpler, having only one sector and no capital. As I explain later, there is no need to introduce capital in the model, as polarisation can be evaluated directly by its effect on relative productivities.

Naturally, this setting brings this model closer to the traditional sorting literature in labour economics, where comparative advantage takes central stage. The pioneer of this literature is undoubtedly in Roy (1951), although he does not provide a mathematical framework to characterise occupational sorting. A well-known formalisation of Roy's framework is in Borjas (1987), used to understand which types of workers migrate from Mexico to the US. Yet, this is an example of models which do not take a general equilibrium approach, in the sense that wages are exogenous to the model, and the corresponding sorting does not alter their values in equilibrium. A closer model is perhaps Sattinger (1975), a paper that seem to have been left out of the references in recent decades. The goal of this paper is to show the conditions under which the distribution of abilities and earnings are different. Perhaps unexpectedly, whenever there is comparative advantage (i.e. workers are different), that is the case. The framework used is of an economy with a continuum of workers who can perform a continuum of tasks. Given comparative advantage, the paper characterises in simple terms the equilibrium conditions in the labour market, conditions that resemble those found here for the case of three tasks. Thus, in a sense the model in this chapter draws from both Autor and Dorn (2013) and Sattinger (1975), adding an empirically relevant set of assumptions in terms of distributions and full mobility of workers, leading to a new model that is very appropriate for the study of job polarisation.

2.3 Model

2.3.1 Production technology

Consider a closed economy with a population of workers $L = 1$. There are three "tasks", or occupations, where workers can be employed.[4] These tasks are: a *manual* task (m),

[4] A task is commonly defined as an indivisible process of production. For example, if the good is a pair of shoes, tasks could be (i) cutting its components, (ii) assembling them, and (iii) packaging. Or,

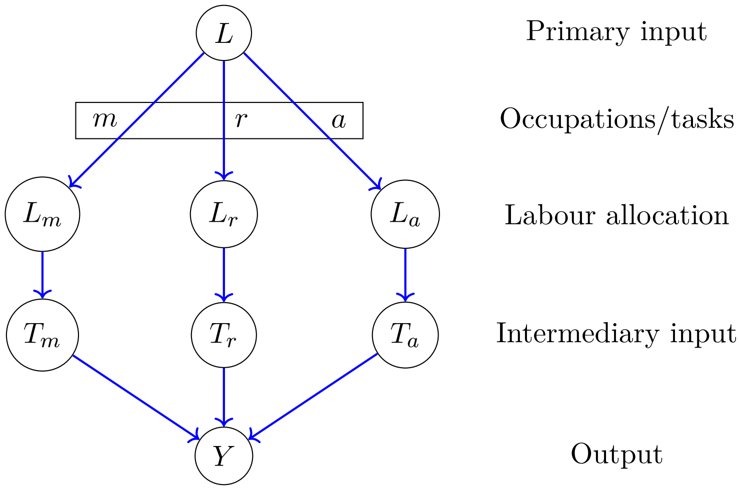

Figure 2.2: Model's structure

a *routine* task (r), and an *abstract* task (a). The output from these tasks is itself an input used to produce one final good, Y. More precisely, let us index tasks by j, and denote the output resulting from workers allocated into task j as T_j. The technology to produce Y is given by the following aggregate production function:

$$Y = \left(\alpha_m T_m^{\phi} + \alpha_r T_r^{\phi} + \alpha_a T_a^{\phi}\right)^{\frac{1}{\phi}} \tag{2.1}$$

where the elasticity of substitution between tasks is $\sigma_Y = \frac{1}{1-\phi}$.[5]

The value of T_j does not depend only on the number of workers allocated into task j but on their ability for that task. As described later, workers have heterogeneous ability for these tasks, making sorting into tasks the key problem to be solved. Naturally, the type of heterogeneity assumed becomes central to the model, so considerable attention is dedicated to these assumptions.

This two-layered framework is summarised in Figure 2.2, with T_j being an intermediate input, and Y being the final good. L is the fixed supply of labour and L_j is the amount of labour allocated into task j, where $L_m + L_r + L_a = 1$.

more broadly, tasks could be (i) design, (ii) manufacture, and (iii) distribution. By assuming that an employed worker contributes only to one task, occupations (i.e employment in a firm) and tasks become synonymous. For more details about the task approach, see Section 3.6 in Chapter 3; also see Autor (2013).

[5]The decision to define production in terms of ϕ instead of σ_Y is simple. One of the key parameters of the labour framework refers to the standard deviation of the log-normal distribution, which standard notation is σ. Using any other symbol for the latter seems much more alien that defining the CES production function in terms of ϕ.

There are four markets in this economy, all perfectly competitive. These are the market for final goods, and the three markets for intermediate inputs T_j.[6] Since there is one final good, all prices and wages are defined in terms of the price of the final good, which acts as numéraire.

2.3.2 Labour framework

There is a continuum of workers of mass 1, indexed by i, where $i \in [0, 1]$. Each worker has a certain ability endowment for performing each of this economy's tasks. These endowments are not defined in terms of a common unit (like IQ) but directly as output per unit of time (e.g. 5 units of "manual task output" per hour).

Ability endowments

Workers' ability endowments are given by the following **ability function**:

$$\eta_j(i) = exp\left[\mu_j + \sigma_j\sqrt{2}\,\mathrm{erf}^{-1}(2i-1)\right] \tag{2.2}$$

where $\eta_j(i)$ denotes the ability of worker i for task j, $\mathrm{erf}(\cdot)$ is the complementary error function, and μ_j and σ_j are parameters. An example of this function is shown in Panel (a) of Figure 2.3. It is monotonic and increasing over i. For example, worker $i = 0.3$ is less productive than worker $i = 0.5$ in task j. This function is chosen because it leads to a log-normal **distribution** of ability for each task.[7] The mean and variance of this log-normal distribution is μ_j and σ_j^2 respectively. In other words, $\eta_j \sim logN(\mu_j, \sigma_j)$. Panel (b) of Figure 2.3 shows the *pdf* of the corresponding distribution. It's algebraic expression is:

$$f(\eta_j) = \frac{1}{\eta_j\sigma_j\sqrt{2\pi}}exp\left[-\frac{(\ln\eta_j - \mu_j)^2}{2\sigma_j^2}\right] \tag{2.3}$$

[6] Strictly speaking, we could also add the market for labour, where workers allocate to firms who produce these intermediate inputs. Yet, through this chapter I assume perfect competition in all markets. As such, these intermediate firms provide no insight into the functioning of the model.

[7] In fact, the inverse of this function is $i(\eta_j) = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{\ln(\eta_j) - \mu_j}{\sigma_j\sqrt{2}}\right)\right]$, which is exactly the cumulative distribution function of a log-normal distribution, where $i \in [0, 1]$. By means of the Inverse Transform Method, it follows that $\eta_j(i)$ distributes log-normal. See Appendix C for the details.

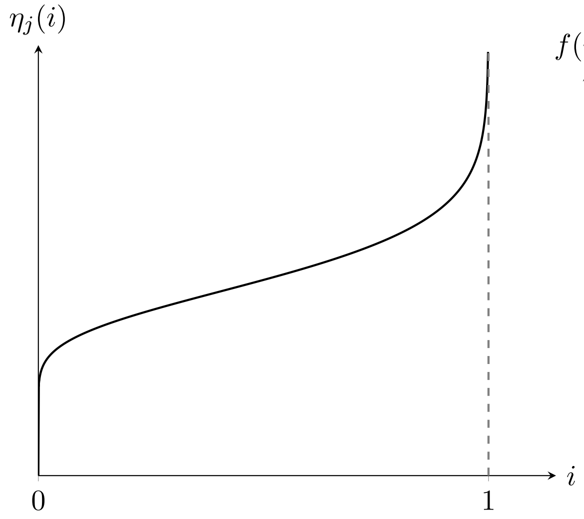

(a) Function defining workers' ability for task j

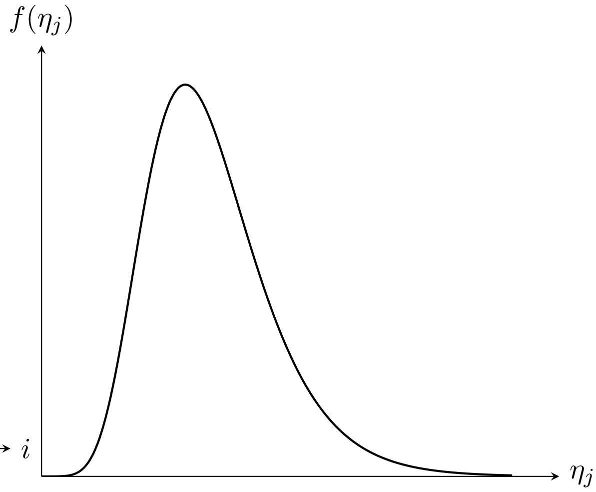

(b) Distribution of ability for task j

Figure 2.3: Ability framework of the model

Since the same function – albeit with different parameters – defines ability of workers among tasks, equation (2.2) implies **absolute advantage** among workers for all tasks. This is, worker i_2 has higher ability than worker i_1 for **every** task, provided that $i_1 < i_2$. This feature is sometimes called hierarchical ability in the labour economics literature.

This assumption is surely very strong and, strictly speaking, unrealistic. In fact, both external research from psychology and educational sciences and the empirical analysis in Chapter 3 rejects it. However, among all the set of assumptions which gives tractability to the model, absolute advantage is the best approximation to the empirical evidence. In effect, what is needed for tractability is to order workers' ability in all tasks use the **same** index i. Conversely, the most general case would be one where there are three indexes, one for each task, where there is total freedom on how workers ability compare between tasks. This would introduce not only further notation but would deny any characterisation of the labour market allocation. The case of absolute advantage is one where the correlation between these indexes is 1. As such, there is no need for three indexes but just one – i.

There are other possible orderings which are also tractable. For instance, we could reverse the ability order for both manual and routine tasks. This is, we could assume that lower i represents strictly *higher* ability than high i for manual and routine tasks, and that lower i represents strictly *lower* ability than high i for abstract tasks. In terms of the three indexes mentioned above, the correlation between the manual and routine indexes is 1, whereas the correlation between both the manual and routine indexes and the abstract index is -1. In this example, perfect correlation means we can reduce the three indexes to just one. In itself, this would lead to no changes in the model or its solution, albeit it would affect the polarisation analysis.

The key question then is which tractable assumption to choose. Empirical evidence indicates that the correlation between "manual ability" and "cognitive ability" is positive. This is, individuals with higher IQ tend to have better spatial and mechanical aptitudes (like dexterity), and tend to develop motor skills faster (Stoeger and Ziegler, 2010; Smits-Engelsman and Hill, 2012; Mihaela et al., 2013). Thus, faced with the need to choose a particular assumption, absolute advantage is the one to provoke the least degree of controversy.

Workers' preferences

Workers' utility comes only from consumption of Y. Without capital in this economy, workers' only source of income are wages coming from their chosen occupation. Furthermore, all workers have an equal endowment of time, which is indivisible. Thus, the only choice workers face is their occupation. Naturally, under this setting workers sort into the job that provides them the highest wage.

Define the wage received by worker i when employed in task j as $w_j(i)$. As shown later (section 2.4.4), due to constant returns to scale in technology and perfect competition, $w_j(i) = \omega_j \eta_j(i)$, where ω_j is the wage *per ability unit* (or wage rate) in occupation j – common for all workers employed in j. In other words, workers' pay is proportional to their ability.

Now, for a given set of wages rates $\{\omega_m, \omega_r, \omega_a\}$, a worker i (strictly) prefers occupation j over all other occupations $-j$ if the wage obtained in j is (strictly) larger than those he would obtain in $-j$. This is:

$$j \succeq -j \quad iff \quad \omega_j \eta_j(i) \geq \omega_{-j} \eta_{-j}(i) \tag{2.4}$$

It is clear from the above equation that workers' abilities are crucial for determining their preferences over tasks. In effect, *ceteris paribus*, the more able a worker is for a certain task, the more likely is that she prefers it over the other tasks. Thus, workers' choices and ultimately the labour market allocation depends crucially on the ability functions $\eta_j(i)$.

Now, it is clear that absolute advantage by itself cannot determine the labour market allocation. We need further assumptions about **comparative advantage** among workers. Thus, to move forward, assume the following:[8]

Assumption 1

$$\sigma_m < \sigma_r < \sigma_a \tag{2.5}$$

The importance of this assumption is in the result it yields. In particular, consider the ratio of ability functions j and k, $\frac{\eta_j(i)}{\eta_k(i)}$. This can be re-written as:

$$exp\left[(\mu_j - \mu_k) + (\sigma_j - \sigma_k)\sqrt{2}\,\text{erf}^{-1}(2i-1)\right]$$

Now, its derivative with respect to i is:

$$exp\left[(\mu_j - \mu_k) + (\sigma_j - \sigma_k)\sqrt{2}\,\text{erf}^{-1}(2i-1)\right](\sigma_j - \sigma_k)\sqrt{2\pi}\,exp\left[\text{erf}^{-1}(2i-1)^2\right]$$

Since the exponential is a strictly positive function, the only factor that determines the sign of this derivative is the relative value of σ_j. Therefore, given **Assumption 1**, it holds that:

Result 1: comparative advantage

$$\frac{\mathrm{d}\left(\frac{\eta_a(i)}{\eta_r(i)}\right)}{\mathrm{d}i} > 0 \qquad \frac{\mathrm{d}\left(\frac{\eta_r(i)}{\eta_m(i)}\right)}{\mathrm{d}i} > 0 \tag{2.6}$$

[8]Chapter 3 finds some support for this assumption. In particular, for all estimations, $\sigma_m < \sigma_a$. As stated there, the estimation of routine ability – and hence σ_r – is more volatile, leading to departures from this assumption in four out of six estimates. For more details, see Table 4.2 later.

These expressions refer to the rate of change of relative abilities over i. It means that as we move from low i to high i, ability increases proportionally more in routine (abstract) relative to manual (routine) occupations. This is, workers closer to $i = 1$ have comparative advantage in producing abstract (routine) tasks compared to routine (manual) tasks, whereas workers closer to $i = 0$ have comparative advantage in producing routine (manual) tasks compared to abstract (routine) tasks.

The conditions in equation (2.6) are quite useful to characterise the equilibrium in the labour market, by means of comparing workers preferences over tasks. In particular, we can search for an indifferent worker between a pair of tasks, to then find the preferences of the remaining workers using **Result 1**. This procedure is enough to map all preferences for tasks among workers, and characterise the final allocation.

1. Manual task *versus* routine task:

   Let us denote the worker that is indifferent between working in occupations m and r as i_1, with $i_1 \in [0, 1]$. Then, worker i_1 is indifferent *iff*, for a given ω_m and ω_r, it holds that:

   $$\omega_m \eta_m(i_1) = \omega_r \eta_r(i_1) \tag{2.7}$$

   Or:

   $$\frac{\eta_r(i_1)}{\eta_m(i_1)} = \frac{\omega_m}{w_r} \tag{2.7'}$$

   The left-hand side of Equation (2.7') is a ratio of the respective ability functions, which vary over i, whereas the right-hand side is (for now) a positive constant. These two functions are shown in Figure 2.4. As mentioned above, the condition $\sigma_m < \sigma_r$ ensures that the ability ratio grows over i. In fact, it can be shown that this ratio is a continuous, strictly monotonic, and increasing function. It follows that the ability ratio crosses the non-negative wage ratio once, at i_1. In consequence, i_1 splits the i interval in two, with workers having the following preferences:

   $$\begin{aligned} m \succ r \quad &\text{for} \quad 0 \leq i < i_1 \\ r \succ m \quad &\text{for} \quad i_1 < i \leq 1 \end{aligned} \tag{2.8}$$

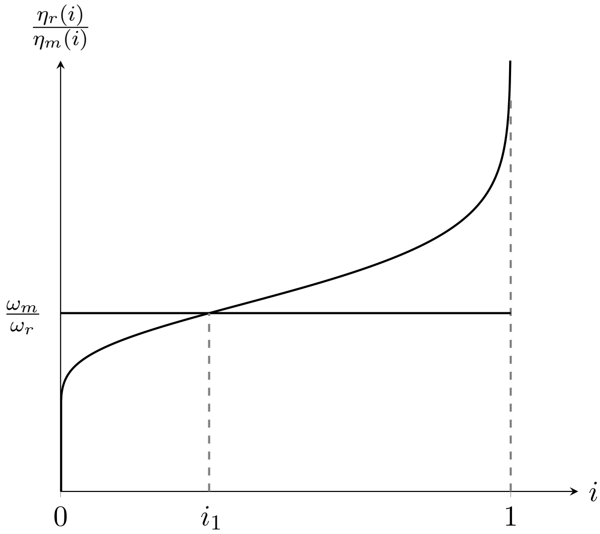

Figure 2.4: Existence and uniqueness of an indifferent worker between m and r tasks

But, does this worker exist? Note that the ability ratio is defined over the $\mathbb{R}^+$ codomain. This means that, for *any* set of positive wages, there is always a worker indifferent among these tasks (i.e. **existence** is satisfied). Moreover, strict monotonicity of this ratio implies that there is only one crossing point between the curves in Figure 2.4. In other words, the indifferent worker is **unique**.

2. Routine task *versus* abstract task:

The procedure for this case is just as above. Denote the indifferent worker between these to tasks as i_2, with $i_2 \in [0, 1]$. Mathematically, for a given ω_r and ω_a:

$$\omega_r \eta_r(i_2) = \omega_a \eta_a(i_2) \tag{2.9}$$

Or:

$$\frac{\eta_a(i_2)}{\eta_r(i_2)} = \frac{\omega_r}{w_a} \tag{2.9'}$$

The graphical representation of this relationship follows just like that in Figure 2.4, and thus it is omitted. Again, due to **Assumption 1**, there is only one crossing point, a consequence of the ability ratio being continuous, strictly monotonic, and increasing, with codomain $\mathbb{R}^+$. Hence, the indifferent worker exists and is unique. In this case, the cut-off point i_2 divides the i dimension in the following way:

$$\begin{aligned} r \succ a \quad &\text{for} \quad 0 \leq i < i_2 \\ a \succ r \quad &\text{for} \quad i_2 < i \leq 1 \end{aligned} \tag{2.10}$$

Table 2.1: Workers preferences over tasks

	i_1	i_2	i_3
$i \in [0, i_n)$	$m \succ r$	$r \succ a$	$m \succ a$
$i \in (i_n, 1]$	$r \succ m$	$a \succ r$	$a \succ m$

3. Manual task *versus* abstract task:

   Similarly, there exists a unique indifferent worker i_3 that is indifferent between occupations m and a. In this case, workers' preferences are:

$$\begin{aligned} m \succ a \quad &\text{for} \quad 0 \leq i < i_3 \\ a \succ m \quad &\text{for} \quad i_3 < i \leq 1 \end{aligned} \tag{2.11}$$

Table 2.1 summarises these results so far, where i_n refers to the three cut-off points defined above. Given these preferences, where do workers allocate? The answer depends on the relative value of i_1, i_2, and i_3. There are six possible orderings for these points. Yet, only one can occur in equilibrium, namely $i_1 < i_3 < i_2$. The other five scenarios are either inconsistent with workers preferences or lead to a task not being performed (impossible in the general equilibrium if $\phi < 1$). In what follows, I only focus on the aforementioned, consistent ordering. For a characterisation of the inconsistent scenarios, see Appendix D.

Labour allocation

Panel (a) in Figure 2.5 presents the resulting labour market allocation when $i_1 < i_3 < i_2$ holds. The horizontal line represents i, which is divided into four intervals. For each, the pairwise comparison of preferences is annotated according to the information in Table 2.1. These comparisons lead to one occupation being favoured by those workers in that interval – shown on top of the line. The terms L_j represent these allocations. More precisely:

$$L_m = i_1 \qquad L_r = i_2 - i_1 \qquad L_a = 1 - i_2 \tag{2.12}$$

where $L_m + L_r + L_a = 1$. Panel (b) in the same figure summarises this result.

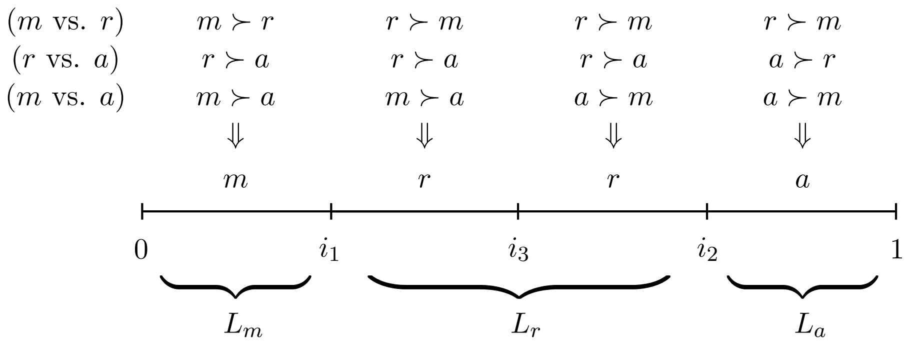

(a) Labour allocation for $i_1 < i_3 < i_2$

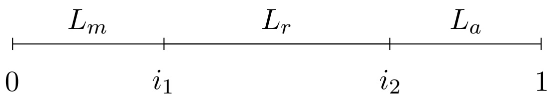

(b) Summary of equilibrium

Figure 2.5: Only consistent labour market allocation in equilibrium

Clearly, in an economy with three tasks, the subdivision of i into three employment intervals, shown in Panel (b) of Figure 2.5, is the simplest allocation possible. This result emerges directly from comparative advantage (**Assumption 1**). In particular, the strictness ($<$) in the relation between σ_js means there is a unique indifference worker between each pair of tasks. This uniqueness is desired if one wants to produce a market allocation with only three intervals along i (this is, with two cut-off points, i_1 and i_2). Removing this strictness means more than two cut-off points along i. The computation of L_i and of T_i would be much more complex, making the characterisation of the solution of the model impossible.[9] But more importantly, allowing more cut-off points introduces need for extra decisions to make regarding the number of intervals, which is not necessarily less arbitrary, nor more informative.

Assumption 1 also involves a second choice, which is the specific sign or order chosen between σ_js ($m < r < a$). This order defines which workers allocate to each tasks. In particular, it implies that those working in abstract tasks are the **best** for abstract tasks among all the population, whereas those working in manual tasks are the **worst** for manual tasks among the population (with those in routine jobs in between). An

[9]In particular, T_i would be defined as the sum of several integrals along i, instead of just one integral, as Section 2.3.2 later shows. The solution of the model – a 2x2 system – becomes unnecessarily cumbersome.

alternative allocation possible is given by $a < r < m$, which means that those working in manual tasks are the best among all the population, whereas those working in abstract tasks are the worst among the population. Such allocation is clearly unsatisfactory. In fact, from all possible combinations, the allocation implicit in **Assumption 1** ($m < r < a$) is by far the most satisfactory empirically. On the one hand, this is very much supported by the empirical analysis in Chapter 3. On the other hand, it is rather natural to conceptualise educational choices (in particular non-compulsory education like university degrees) as a sorting mechanism by which the most able individuals among a population choose to become high-skilled (i.e. those with comparative advantage in cognitive abilities). In effect, an higher-education degree is an mechanism by which individuals acquire credentials "enabling" them to perform high-skilled jobs. Importantly, these credentials still allow them to work in unskilled jobs (e.g. as Uber drivers). Yet, the reverse is generally not true; low-educated workers are many times barred from high-skilled jobs by institutional constraints, regardless of how good they might be on such jobs. Therefore, an allocation where the best among a given population sort into cognitive-intensive jobs is certainly a desirable feature. An associated cost of this, imposed by tractability requirements, is the result that manual workers are the worst for manual tasks among the population. This is clearly a simplification, but the best among the alternatives available.

To summarise, whilst acknowledged that the sorting resulting from comparative advantage is, in its strict form, unrealistic, among all simplifications possible it is the more plausible one – and partly justified by the data. Thus, just as in the case of absolute advantage, both tractability and empirical relevance justifies **Assumption 1**.

The following definition summarises the result so far:

Definition 1 (Optimal labour market allocation)**.** *Denote L the set of employment allocation $\{L_m, L_r, L_a\}$, and denote W the set of relative wages $\{\frac{\omega_m}{\omega_r}, \frac{\omega_r}{\omega_a}\}$. Then, the pair $\{L, W\}$ is an optimal labour market (OLM) allocation iff:*

a) for each task j:

$$\omega_j \eta_j(i) \geq \omega_{-j} \eta_{-j}(i) \quad \forall i \in L_j$$

b) $i_1 < i_3 < i_2$

The first condition ensures that, for a given a set of wages, workers are better off at their chosen occupation. The second condition ensures that all occupations are performed (again, see Appendix D). Notice that **there are infinite OLM allocations**, due to the continuous nature of the ability functions and the full flexibility of wages.[10] As shown later, only one of them occurs in the general equilibrium.

A bit of intuition

It is useful to bring further intuition to the surface by relating the ability functions, **Assumption 1**, and the final allocation. Consider a worker's potential earnings in a given occupation: $w_j(i) = \omega_j \eta_j(i)$. Because wage rates are always positive if $\phi < 1$, this earnings function inherits the properties of $\eta_j(i)$. In particular: (i) its range is always between $-\infty$ and $+\infty$, by virtue of $\eta_j(0) = -\infty$ and $\eta_j(1) = +\infty$ respectively; (ii) this range is independent of the values of μ_j and σ_j (except in the trivial degenerate case of $\sigma_j = 0$, where workers are identical); (iii) its slope is always positive; and (iv) this slope depends positively on the wage rate, μ_j, and σ_j. Since the range of the function is fixed, the slope relates to the degree of "compression" of the function.

Figure 2.6 shows the labour market equilibrium from the perspective of these earnings functions, under **Assumption 1**. As demonstrated earlier, this assumption implies comparative advantage among workers and tasks, leading to three cut-off points, defined by the crossing of pairs of ability functions. These points represent the indifferent workers found before. Each worker i compares the values of the three lines (i.e. her total pay on each occupation), choosing the highest. As it is clear from the graph, the allocation is defined entirely by i_1 and i_2. Importantly, this graph is based on the **equilibrium** wage rates. Recall that only one ordering is consistent with the general equilibrium. To achieve it, relative wages will adjust, until the allocation looks like that in Figure 2.6.

[10] If not intuitive yet, consider a particular OLM allocation $\{L, W\}_A$, where – by definition – $i_1 < i_3 < i_2$ holds. From Figure 2.3 and its equivalent for other tasks, it is clear that a small increase in w_r reduces i_1 and increases i_2, while i_3 remains constant. Since the inequality $i_1 < i_3 < i_2$ still holds, the new allocation $\{L, W\}_B$ is necessarily an OLM allocation too. The continuous nature of the variables of the model leads to the infinite cardinality of the OLM set.

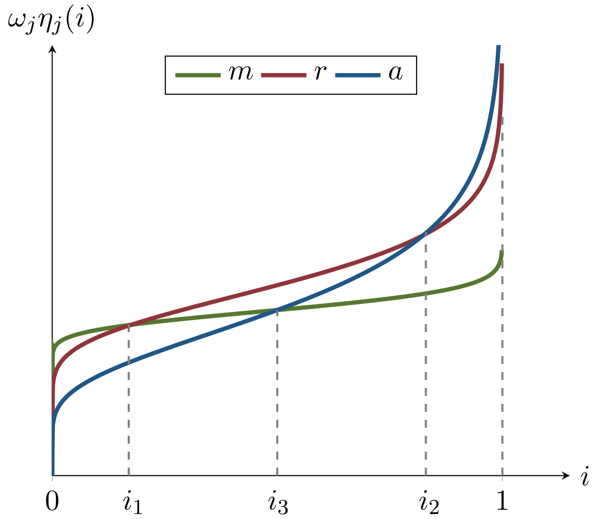

Figure 2.6: Labour market allocation under $\sigma_m < \sigma_r < \sigma_a$ (Assumption 1)

Now, what if instead of **Assumption 1**, the σ_js are equal? It is trivial to show that in this scenario, $\frac{\mathrm{d}\left(\frac{\eta_j(i)}{\eta_{-j}(i)}\right)}{\mathrm{d}i} = 0$. This is, **there is no comparative advantage** among workers. Naturally, there is still absolute advantage – and heterogeneity. But this is not enough to ensure a definite sorting allocation. In fact, the equilibrium is undefined in the sense that *who works where* is irrelevant. In terms of Figure 2.6, the three lines would coincide.

The role of μ_j in the equilibrium allocation is more difficult to evaluate. It was already shown that μ_j has no effect on comparative advantage. But surely, if the level of ability of workers for just one task is altered, wouldn't the sorting change? In effect, a *ceteris paribus* increase in μ_j means workers are more productive in this task. This changes the demand for tasks, thereby altering the wage rates, which in this exercise were considered exogenous.[11] This is, a complete analysis of the role of μ_j requires a general equilibrium approach, which has not been introduced yet. There is however one case which is intuitive enough. As said before, if σ_js are all equal, there is no comparative advantage, and workers are indifferent between occupations. Now, imagine μ_a increases. This means everyone becomes more able at the abstract task. However, comparative advantage still does not exist. Clearly, if wage rates do not change, every worker would be better off in the abstract occupation. But this is not an equilibrium. This excess

[11]The particular direction of change depends on the elasticity of substitution between tasks. If they are gross complements ($\phi < 0$), firms demand less of j and more of the less productive tasks. If they are gross substitutes ($\phi > 0$), the converse is true. For more details, see Section 2.5.

supply of abstract workers would change the relative wages in detriment of ω_a, until all excess supply is eroded, and the economy is back to the original indifference state. This shows an example where μ_j affects wages and the allocation of workers, without altering comparative advantage; the original sorting order $(m - r - a)$ holds, and $i_1 < i_2$. If the change in μ_j occurs when **Assumption 1** holds, the allocation might still change, but not the sorting order. Again, wage rates ensure that the only equilibrium possible in this economy $(i_1 < i_2)$ is achieved, and workers are employed in every task.

Output from tasks

Now that the set of OLM allocations has been characterised, it is possible to derive the output from tasks, T_j, which are then inputs in the production of the final good. Each worker's contribution to a task depends directly on his ability level. For simplicity, I assume here that T_j is given by the **sum** of the ability for task j of those workers employed on it. This is:

$$T_j \equiv \int_{i \in L_j} \eta_j(i)\, \mathrm{d}i \tag{2.13}$$

Given the result in Equation (2.12), these outputs are:

$$T_m \equiv \int_0^{i_1} \eta_m(i)\, \mathrm{d}i \qquad T_r \equiv \int_{i_1}^{i_2} \eta_r(i)\, \mathrm{d}i \qquad T_a \equiv \int_{i_2}^{1} \eta_a(i)\, \mathrm{d}i \tag{2.14}$$

Using the particular functional form for ability given in Equation (2.2), each task's output is:

$$\begin{aligned}
T_m &= \Lambda_m \left[1 - \mathrm{erf}\left(\frac{\sigma_m}{\sqrt{2}} - \mathrm{erf}^{-1}\left(2i_1 - 1\right)\right)\right] \\
T_r &= \Lambda_r \left[\mathrm{erf}\left(\frac{\sigma_r}{\sqrt{2}} - \mathrm{erf}^{-1}\left(2i_1 - 1\right)\right) - \mathrm{erf}\left(\frac{\sigma_r}{\sqrt{2}} - \mathrm{erf}^{-1}\left(2i_2 - 1\right)\right)\right] \\
T_a &= \Lambda_a \left[1 + \mathrm{erf}\left(\frac{\sigma_a}{\sqrt{2}} - \mathrm{erf}^{-1}\left(2i_2 - 1\right)\right)\right]
\end{aligned} \tag{2.15}$$

where $\Lambda_j = \frac{1}{2} exp\left(\mu_j + \frac{\sigma_j^2}{2}\right)$. Figure 2.7 shows these outputs graphically, for some arbitrary parametrisation. They correspond to the area below the function within the relevant range, as per their definition in Equation 2.13. Naturally, the expansion of employment into a task leads to more output from it, and vice-versa.

Finally, note that functions in Equation 2.15 do not seem like a set of supplies since they do not involve wages. Instead, they are functions of the cut-off points, i_1 and i_2.

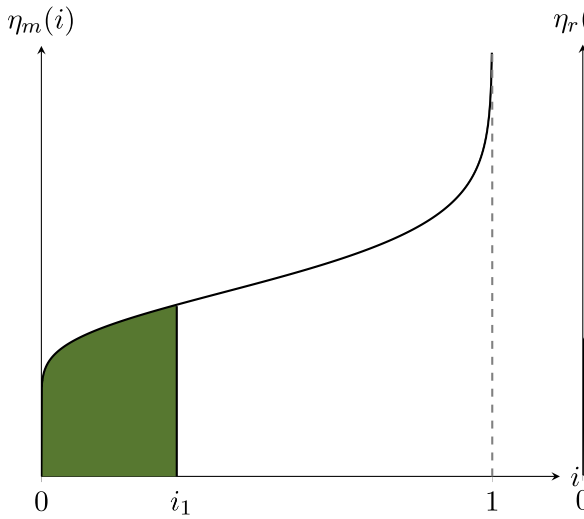

(a) $T_m \equiv \int_0^{i_1} \eta_m(i)\,\mathrm{d}i$

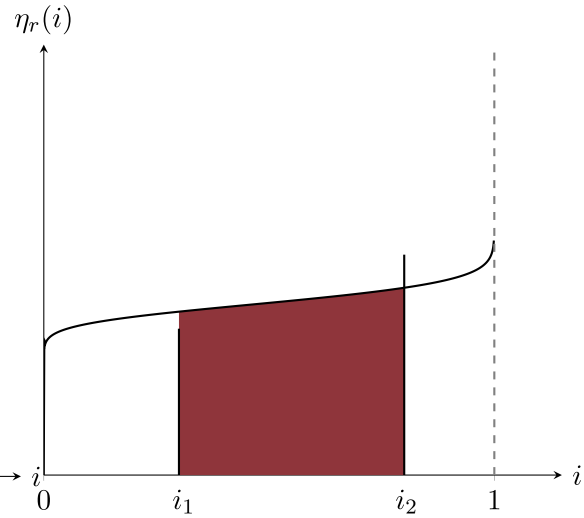

(b) $T_r \equiv \int_{i_1}^{i_2} \eta_r(i)\,\mathrm{d}i$

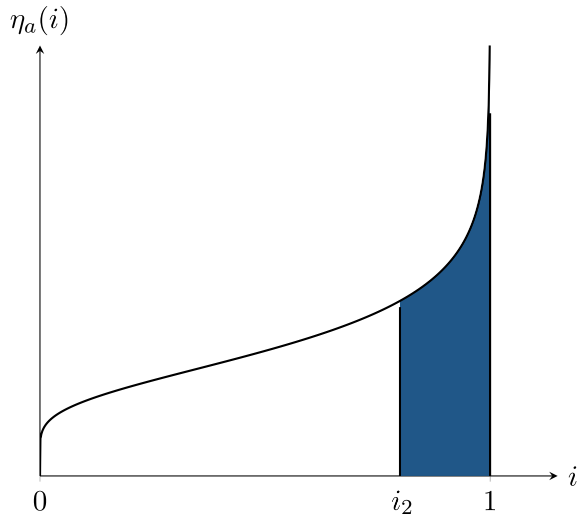

(c) $T_a \equiv \int_{i_2}^{1} \eta_a(i)\,\mathrm{d}i$

Figure 2.7: Example of tasks output with $\sigma_m < \sigma_r < \sigma_a$

Yet, recall these points depend directly on the relative level of wage rates, as equation (2.7) and (2.9) indicate. Therefore, functions in Equation 2.15 do represent an implicit set of supplies of tasks.

2.4 Solution[12]

Definition 1 describes the infinite set of labour market equilibria (in terms of employment and wages) that are consistent with both workers' preferences **and** positive employment in all tasks. Yet, different labour allocations result in different output levels. From a benevolent social planner perspective, the interest is with that or those allocations which equilibrate the labour market **and** maximises output. Let us define these allocations as follows:

Definition 2 (Optimal global (0G) allocation)**.** *If Y denotes output in this economy, then the set $\{L, W, Y\}$ is an optimal global (OG) allocation iff the pair $\{L, W\}$ is an OLM allocation which maximises Y over all possible OLM allocations.*

How can we find this global equilibrium? From the perspective of markets, the equilibrium is achieved by competitive forces that adjust relative wage rates in order to induce workers to allocate to tasks, leading to a combination of inputs T_j that maximises aggregate production. For example, consider an initial OLM, where workers are allocated to tasks, and $i_1 < i_2$, given a set of relative wage rates. Now, say that firms can increase their profits by using less T_r and more T_a. This gives the incentive to (all) firms to raise the relative wage for abstract tasks, inducing workers to reallocate. This happens until the profit gains are exhausted. This is the OG allocation.

Instead of solving the model from the perspective of firms, workers, and markets, the solution is characterised from the perspective of a benevolent social planner. The first welfare theorem ensures they are the same – something it is shown below. This theorem can be applied here because all its requirements hold, namely competitive markets, perfect information, no transaction costs, and no externalities. The last part of this section reverts back into firms and markets.

[12]This section presents the solution to the model outlined in the previous section, i.e. with three tasks and log-normal ability distributions. The solution of a generalised model in terms of ability distributions and number of tasks is presented in Appendix E.

2.4.1 Social planner's maximisation problem

Consider a social planner interested in maximising final output. This output is a function of the three type of tasks, T_j. Yet, as equation (2.15) shows, these are functions of two variables: i_1 and i_2. In consequence, it is enough to find these two endogenous variables, from which the complete solution of the model can be found. In particular, the maximisation problem is:

Maximisation problem

$$\max_{\{i_1,i_2\}} \quad Y$$

subject to:

$$Y = \left(\alpha_m T_m^{\phi} + \alpha_r T_r^{\phi} + \alpha_a T_a^{\phi}\right)^{\frac{1}{\phi}}$$
$$T_m = \Lambda_m \left[1 - \text{erf}\left(\frac{\sigma_m}{\sqrt{2}} - \text{erf}^{-1}\left(2i_1 - 1\right)\right)\right]$$
$$T_r = \Lambda_r \left[\text{erf}\left(\frac{\sigma_r}{\sqrt{2}} - \text{erf}^{-1}\left(2i_1 - 1\right)\right) - \text{erf}\left(\frac{\sigma_r}{\sqrt{2}} - \text{erf}^{-1}\left(2i_2 - 1\right)\right)\right]$$
$$T_a = \Lambda_a \left[1 + \text{erf}\left(\frac{\sigma_a}{\sqrt{2}} - \text{erf}^{-1}\left(2i_2 - 1\right)\right)\right]$$
$$0 < i_1 < i_2 < 1$$

where $\Lambda_j = \frac{1}{2} exp\left(\mu_j + \frac{\sigma_j^2}{2}\right)$.[13] The last constraint ensures the solution is an Optimal Labour Market allocation.

Here it is also interesting to summarise the parameter space. First, the distributions of ability assume $\mu_j > 0$ and $\sigma_j > 0$, $\forall j$. Second, an interior solution (i.e. all tasks are used on Y) is assured by $-\infty < \phi < 1$. Third, **Assumption 1** requires $\sigma_m < \sigma_r < \sigma_a$. Finally, $\alpha_j > 0$, $\forall j$.

[13] It is reminded to the reader that σ_j refers to the parameter of the ability distributions, different from σ_Y, which is the elasticity of substitution. The latter is implicit in the parameter ϕ, and equivalent to $\sigma_Y = \frac{1}{1-\phi}$

2.4.2 Existence and uniqueness of Optimal Global allocation

Let us denote the aforementioned objective function as $f(i_1, i_2)$ and denote the support set $0 < i_1 < i_2 < 1$ by Ω. Then, $f : \Omega \to \Re$. Note that f is a continuous function on Ω. Furthermore, Ω is bounded and **open**, as it does not include any boundary (e.g. $i_1 = i_2$ is not part of the set). These boundaries entail zero supply of a task, hence they are not part of the possible OLM allocations in equilibrium. Nonetheless, it is useful for now to consider the **closed** set Ω^* given by $0 \leq i_1 \leq i_2 \leq 1$. Under these conditions, the Extreme Value Theorem assures the existence of at least one maximum of f over Ω^*. Now, under the assumption of imperfect substitutions ($\phi < 1$), positive demand for tasks rules out the optimality of any allocation with zero employment in a task. In other words, provided $\phi < 1$, the maximum of f is always inside Ω^*, i.e. in Ω. In conclusion, the solution of the above maximisation problem **exists**.

The objective function $Y = f(i_1, i_2)$ can be proved to be strictly concave over Ω.[14] Given this property and the fact that Ω is a convex set, the maximisation problem posed above is a "convex optimisation problem". It is a corollary from this category of optimisation problems that any local optimum is a global optimum. Furthermore, this optimum is **unique**, because f's convexity is strict. Finally, the uniqueness of the equilibrium also implies its stability, in the case of any exogenous change in the model.

2.4.3 Optimal Global allocation

Having proved that the solution of the maximisation problem exists and is unique, we can move further. The respective FOCs from the problem yield the following equalities:

$$\left(\frac{T_r}{T_m}\right)^{1-\phi} = \left(\frac{\alpha_r}{\alpha_m}\right)\left(-\frac{\frac{\partial T_r}{\partial i_1}}{\frac{\partial T_m}{\partial i_1}}\right) \tag{2.16}$$

$$\left(\frac{T_a}{T_r}\right)^{1-\phi} = \left(\frac{\alpha_a}{\alpha_r}\right)\left(-\frac{\frac{\partial T_a}{\partial i_2}}{\frac{\partial T_r}{\partial i_2}}\right) \tag{2.17}$$

[14] Sadly, an algebraic proof of this is not possible, given the complexity of the function in question (in particular the use of the error function). Using specialised computing languages like Python or Julia cannot provide the proof either. Yet, a graphical representation of this function "proves" the function is concave. See Figure 2.8 below.

Plugging in the derivatives of equation (2.15), and combining exponential terms results in:[15]

2x2 System

$$\left(\frac{T_r}{T_m}\right)^{1-\phi} = \left(\frac{\alpha_r}{\alpha_m}\right) exp\left[(\mu_r - \mu_m) + (\sigma_r - \sigma_m)\sqrt{2}\,\text{erf}^{-1}\left(2i_1^* - 1\right)\right] \tag{2.18}$$

$$\left(\frac{T_a}{T_r}\right)^{1-\phi} = \left(\frac{\alpha_a}{\alpha_r}\right) exp\left[(\mu_a - \mu_r) + (\sigma_a - \sigma_r)\sqrt{2}\,\text{erf}^{-1}\left(2i_2^* - 1\right)\right] \tag{2.19}$$

This is a 2x2 non-linear system, with endogenous variables i_1 and i_2, and exogenous variables α_j, μ_j, σ_j, and ϕ. Unfortunately, this system yields no closed form solutions for i_1^* and i_2^*.[16] It is however possible to use the Implicit Function Theorem and the Cramer's Rule in order to fully characterise the derivative of these two solutions with respect to all the model's parameters. This is shown in Section 2.5.

Despite the above limitation, the model can be solved numerically and simulated for different parametrisations. Figure 2.8 shows an example for a given parametrisation.[17] The optimal values are $i_1^* = 0.32$ and $i_2^* = 0.89$. Panel (a) confirms the strict concavity of this function. The black dot in the $\{i_1, i_2\}$ plane corresponds to the OG allocation. Panel (b) presents the isoquants of Y in a heat map. The darker the colour the higher the value of f.

2.4.4 Complete solution

The solution to all the endogenous variables of the model derives from the OG allocation $\{i_1^*, i_2^*\}$. In particular, labour and relative wages are given by equations (2.12), (2.7') and (2.9'). Similarly, i_1^* and i_2^* pin down the optimal values of T_j and Y from equations (2.15) and (2.1) respectively. Finally, absolute wage rates ω_j^* come from the competitive labour market condition of wages equalising the marginal product of tasks:

$$\frac{\partial Y}{\partial T_j} = \omega_j \tag{2.20}$$

[15]The * attached to endogenous variables show that these are optimal solution.

[16]There is only case where a closed-form solution exists is in the rather uninteresting (and unrealistic) case of homogeneous ability, i.e. if $\sigma_j = 0$, meaning $\eta_j(i) = exp(\mu_j) = c_j$.

[17]In particular, assumptions are: $\phi = 0.5, \mu_j = 1\,\forall j, \sigma_m = 0.25, \sigma_r = 0.35$ and $\sigma_a = 0.45$.

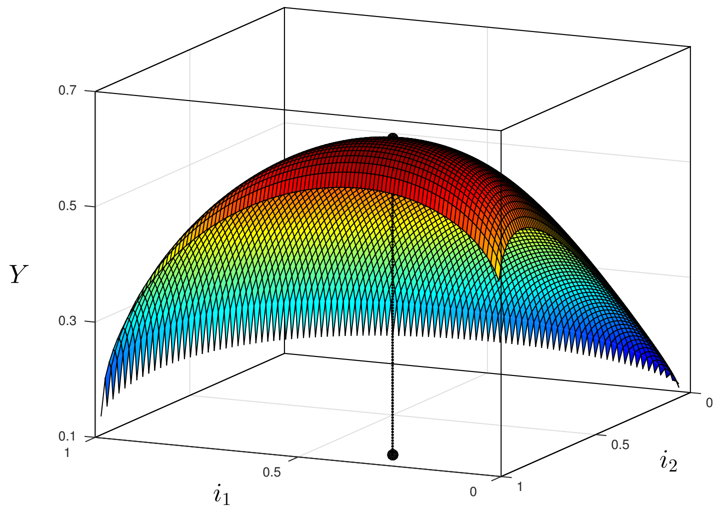

(a) Global maximum of $Y = f(i_1, i_2)$ at (i_1^*, i_2^*)

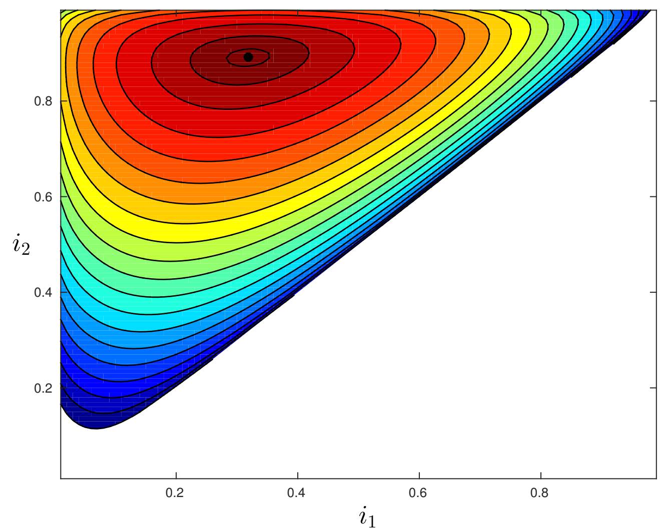

(b) Heat map showing isoquants of Y and global maximum at (i_1^*, i_2^*)

Figure 2.8: Solution for example parametrisation, with $i_1^* = 0.32$ and $i_2^* = 0.89$

From equation (2.1) we know that:

$$\frac{\partial Y}{\partial T_j} = \alpha_j Y^{1-\phi} T_j^{\phi-1} \tag{2.21}$$

Therefore, the equilibrium wage rate for task j is:

$$\omega_j^* = \alpha_j (Y^*)^{1-\phi} (T_j^*)^{\phi-1} \tag{2.22}$$

Note that ω_j^* is the wage *per ability unit*, equal for every worker in occupation j. The income received by worker i, denoted $w^*(i)$, depends on his ability level for the task where he is employed. More precisely, since Y's technology has constant returns to scale and **tasks** markets are competitive, payment to inputs exhaust the final product. This is:

$$Y^* = \omega_m^* T_m^* + \omega_r^* T_r^* + \omega_a^* T_a^* \tag{2.23}$$

Furthermore, recall that T_j is defined as the sum of the ability in task j of workers employed in that occupation:

$$T_j^* \equiv \int_{i \in L_j^*} \eta_j(i) \, \mathrm{d}i \tag{2.24}$$

Consequently, tasks' technology has constant returns to scale too. Now, since the **labour** market is competitive, each of the above payments is exhausted when "paid" to respective workers. This is:

$$\omega_j^* \int_{i \in L_j^*} \eta_j(i) \, \mathrm{d}i = \int_{i \in L_j^*} w^*(i) \, \mathrm{d}i \tag{2.25}$$

The above reduces to:

$$\omega_j^* \eta_j(i) = w^*(i) \quad \forall i \in L_j^* \tag{2.26}$$

This proves that each worker is paid proportionally to his ability. From this set of wages we can study the income distribution and construct inequality indexes.

2.4.5 Markets of tasks

Having characterised the solution of the model from the perspective of the social planner, we can bring the analysis back to the perspective of firms, workers, and the market for inputs T_j. This analysis starts from a central insight arising from the previous subsection, namely, that **the wage rates ω_j correspond to the price of inputs T_j**. In other words, they are the equilibrium price that arises in the market for each tasks. This can be seen in equation (2.23) above, where the payment to tasks fully exhaust the product.

Now, in the model there are three inputs T_j, meaning there are three markets for tasks. As in any general equilibrium model, **the interest is not on each individual input but on relative inputs**. As such, here we are going to focus on the "market" for $\frac{T_m}{T_r}$ and $\frac{T_r}{T_a}$.[18] Naturally, these are interdependent. Additionally, by Walras' Law, finding the equilibrium on the market for tasks automatically yields the solution for the final goods' market.

The demand for tasks comes from firms' maximisation of output. More precisely, if a firm's profit is given by $\Pi = Y - (\omega_m T_m + \omega_r T_r + \omega_a T_a)$, it follows that the relative demands for tasks are:

$$\begin{aligned} \frac{\omega_m}{\omega_r} &= \frac{\alpha_m}{\alpha_r}\left(\frac{T_m}{T_r}\right)^{\phi-1} \\ \frac{\omega_r}{\omega_a} &= \frac{\alpha_r}{\alpha_a}\left(\frac{T_r}{T_a}\right)^{\phi-1} \end{aligned} \tag{2.27}$$

For the case of $\phi < 1$ (i.e. imperfect but positive substitution among tasks), these demands are downward slopping and convex, as Figure 2.9 shows.[19] Similarly, recall that tasks' inputs – equation (2.15) – are functions of the cut-off points i_1 and i_2. This is:

$$\begin{aligned} T_m &= g_m(i_1) \\ T_r &= g_r(i_1, i_2) \\ T_a &= g_a(i_2) \end{aligned} \tag{2.15'}$$

where g_j denotes an arbitrary function. i_1 and i_2 are themselves functions of the relative

[18] The use of the word market is rather loose. Strictly speaking, there are only markets for individual tasks. The "market" for relative tasks is a general equilibrium artificial construct which summarises the relative demand and relative supply of individual tasks, useful to understand general equilibrium changes.

[19] If $\phi = 1$ (i.e. $\sigma_Y = \infty$), these demands are perfectly elastic.

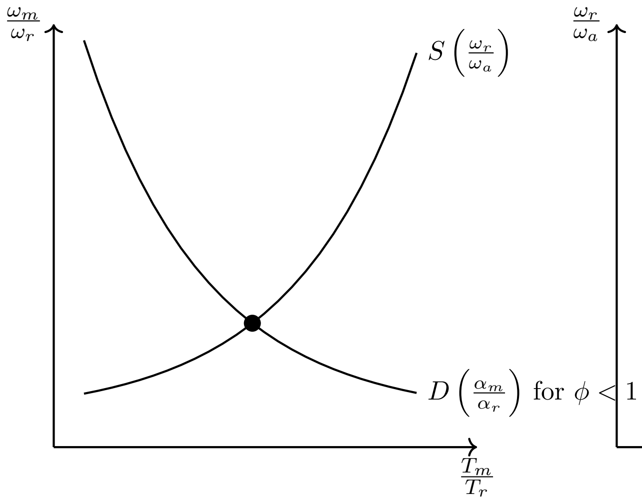

(a) Relative demand and supply of $\frac{T_m}{T_r}$

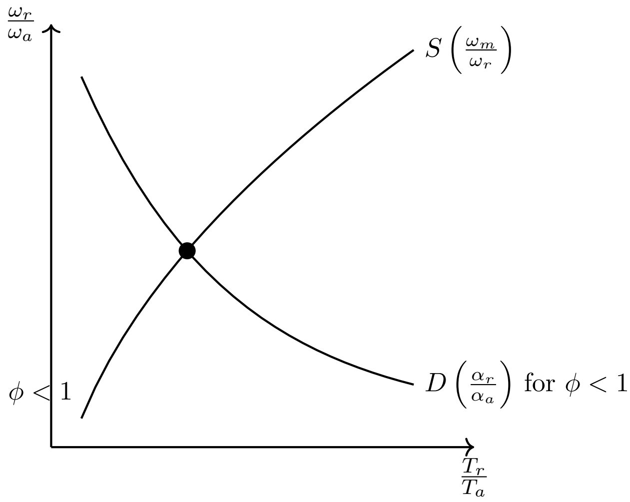

(b) Relative demand and supply of $\frac{T_r}{T_a}$

Figure 2.9: Equilibrium of relative demand and supply for tasks

wages, as equations (2.7') and (2.9') indicates. In particular:

$$\begin{aligned} i_1 &= g_1\left(\frac{\omega_m}{\omega_r}\right) \\ i_2 &= g_2\left(\frac{\omega_r}{\omega_a}\right) \end{aligned} \tag{2.28}$$

where g_n denotes an arbitrary function. Combining equations (2.15') and (2.28) yields the following supplies of tasks:

$$\begin{aligned} \frac{T_m}{T_r} &= S_1\left(\frac{\omega_m}{\omega_r}, \frac{\omega_r}{\omega_a}\right) \\ \frac{T_r}{T_a} &= S_2\left(\frac{\omega_m}{\omega_r}, \frac{\omega_r}{\omega_a}\right) \end{aligned} \tag{2.29}$$

where S_n denotes an arbitrary function. Figure 2.9 shows two arbitrary supplies.[20] Very importantly, they depend not only on their own price but on the other market's price. The simultaneous equilibrium in both markets determines the model's solution.[21]

[20]As already mentioned, there is no algebraic way to write the supply of tasks S_1 and S_2 due to the complexity of η_r and η_a. More precisely, it is impossible to find $g_1(\cdot)$ and $g_2(\cdot)$. Yet, we do know that they are upward slopping because relative wages $\frac{\omega_m}{\omega_r}$ and $\frac{\omega_r}{\omega_a}$ are positively related to i_1 and i_2 respectively, which themselves are positively related to $\frac{T_m}{T_r}$ and $\frac{T_r}{T_a}$. Similarly, it can be shown that these supplies depend negatively on the cross market wages.

[21]Note that describing the model in this way leads to a system of four equations with four unknowns $(\frac{T_m}{T_r}, \frac{T_r}{T_a}, \frac{\omega_m}{\omega_r}, \frac{\omega_r}{\omega_a})$. Yet, this system can be further reduced to a 2x2 system by replacing equation (2.27) into equation (2.29), matching the dimensionality of the problem from the social planner perspective.

2.5 Job polarisation

The relatively simple model presented above can be used to understand job polarisation (JP henceforth). This is defined as the process of labour reallocation from routine tasks into **both** manual and abstract tasks. In the model, JP occurs when L_r falls and L_m and L_a rise; or, in terms of the two equilibrium cut-off points, i_1 increases **and** i_2 falls. There are essentially two independent sources of JP in our framework:[22] (i) an exogenous change in the labour (and tasks) supply due to education, training, learning, etc, reflected in variations in the parameters of the ability distributions (μ_j, σ_j); and (ii) an exogenous technological process that affects the relative demand for tasks (through α_j). How each of these factors lead to JP is now studied, using both a mathematical approach (via partial derivatives and the Implicit Function Theorem) and a graphical analysis. Naturally, complex combinations of these phenomena could also lead to JP, something that it is briefly discussed.

It was mentioned in the literature review that the channels proposed by which JP arises are several. Yet, the existing literature has not yet explored the role workers' ability might play on polarisation. As such, the respective analysis in this section provides a novel contribution to the literature, by adding another channel that might explain polarisation.

2.5.1 Labour supply

According to the model, workers' ability does not vary over time. Yet, in real economies there are many reasons why ability distributions might vary over time. First, even if workers' ability is time-invariant, there are considerable changes in the labour force. Workers leave employment toward unemployment or inactivity, temporarily o permanently, whereas young workers enter the labour force. Strictly speaking, the model here is consistent with entry and exit. Since workers' sorting depends solely on the

[22] There is a third independent source of polarisation, which is a change in the elasticity of substitution. However, in contrast with the two cases presented here, is it not possible to produce a general rule under which job polarisation always occur. The desired result only arises under a particular combination of parametrisations, making this source of polarisation too obscure, and probably not very interesting. See Appendix F for an analysis of ϕ and polarisation.

current periods' wage rates, the solution of the model is independent of the past and future. For a given distribution, the solution of the model would always exist. There is however a limitation. If we want to use this particular log-normal specification of the model to understand changes between two periods, the type of entry/exit must conform to it. This is, the distribution in every period must still be log-normal, albeit with different parameters. A dynamic departure from this setting (e.g. if distributions are no longer log-normal in another period) introduces considerable complexities to the analysis, which might obscure the potential insights of the exercise. As such, the analysis here is still a simplification.

A second source of change in ability – or perhaps more precisely, workers' productivity – is due to training, or learning by doing. The model here is silent about such processes, but it could certainly be introduced. Still, it is possible to think on reasonable changes in ability for the majority of the population. For example, the widespread adoption of computers – ubiquitous know – means that over a given period a great number of workers in the economy improved their productivity at routine tasks. Naturally, the rate of change was not the same for everyone, but that is not required for the distribution of ability to retain its log-normality – any linear transformation of workers' ability does (e.g. $c_0 + c_1\eta_r(i)$, common across workers).

The exercise here is relatively simple. It aims to understand how changes in the ability parameters might lead to polarisation. This requires to evaluate the role of μ_j and σ_j on i_1^* and i_2^*, which ultimate define job polarisation. In particular, we need to characterise the **sign** of

$$\frac{\partial i_1^*}{\partial \mu j} \qquad \frac{\partial i_2^*}{\partial \mu j} \qquad \frac{\partial i_1^*}{\partial \sigma_j} \qquad \frac{\partial i_2^*}{\partial \sigma_j}$$

From here, we can deduce the respective change in employment shares using equation (2.12).

As said earlier, the model does not offer a closed form solution from where these partial derivatives can be deduced. Still, using the Implicit Function Theorem and the Cramer's Rule, these derivatives can be fully characterised. This theorem is useful because it tells us when the aforementioned derivatives exists. In our case, the model is very "regular"

Table 2.2: Sign of $\frac{\partial L_j^*}{\partial \mu_j}$ and $\frac{\partial L_j^*}{\partial \sigma_j}$, for different values of ϕ

Parameters	L_m^*	L_r^* $\phi < 0$	L_a^*	L_m^*	L_r^* $\phi = 0$	L_a^*	L_m^*	L_r^* $\phi > 0$	L_a^*
μ_m	-	?	+	0	0	0	+	?	-
μ_r	+	-	+	0	0	0	-	+	-
μ_a	+	?	-	0	0	0	-	?	+
σ_m	-	?	+	-	?	+	-	?	+
σ_r	+	?	-	+	?	-	+	-	+
σ_a	+	?	-	+	?	-	+	?	-

in the sense that there are no discontinuities on any function, and the Optimal Global allocation is always interior (given $\phi < 1$). Therefore, these derivatives exist all over the parameter space.[23] Additionally, Cramer's Rule provides a simple method to find them. Here I do not provide a step-by-step outline of how these derivatives are found. For details on the theorem and some examples, see Appendix F.

Table 2.2 presents the conclusion from this mathematical exercise, from the perspective of employment shares. The signs on the table represent the direction in which the **equilibrium** employment shares would change, after a *ceteris paribus* variation in a given parameter. A cell containing "?" means the direction of change is ambiguous (i.e. it depends on the combination of multiple parameters). Importantly, some results vary with respect to the sign of ϕ.

The intuition behind changes in μ_j is relatively easy to grasp. A *ceteris paribus* increase in μ_j means all workers are more able at this task (e.g. they produce more per unit of time). This is, $\eta_j(i)$ increases for every i. In consequence, *ceteris paribus*, the output of those workers in j (T_j) increases. Firms react to this depending on the sign of ϕ. If there is relatively high substitution between tasks in production ($\phi > 0$, equivalently to $\sigma_Y > 1$), the production of the final good can be expanded the most by changing the T_j input mix in favour of j. This is achieved by reallocating workers from all other tasks into j. In a competitive market like here, this requires a change in relative wage rates in favour of j. Conversely, there is a contraction in employment in the enhanced task if substitution between them is relatively difficult (i.e. $\phi < 0$, or $\sigma_Y < 1$). In the

[23] Recall that this space is defined by $-\infty < \phi < 1$, $\alpha_j > 0$, $\mu_j > 0$, $\sigma_j > 0$, and **Assumption 1**, equivalent to $\sigma_m < \sigma_r < \sigma_a$.

latter case, firms benefit from this increase in workers' ability by releasing resources (labour) from manual tasks, which reallocate into other tasks. In the pivotal case of Cobb-Douglas ($\phi = 0$), there is no benefit of reallocating employment, as the marginal change in output from these reallocations off-set each other.

Notice that for changes in μ_m and μ_a, there is an ambiguous change in L_r. To understand why, consider the case of an increase in μ_m, when tasks are relatively easy to substitute ($\phi > 0$). As just said, the "own-task" effect is positive, which means workers switch into manual occupations. Now, because of absolute advantage, these workers come directly from routine tasks. Yet, part of this drain on routine workers is compensated by some who switch from abstract to routine. In other words, there is only mobility "downwards". This is necessarily the case because ability in manual tasks increased with respect to **both** routine and abstract, leading to an unambiguous fall in L_a. The same occurs in other cases, where the effect of μ_m or μ_a on L_m and L_a is always defined (see table). In summary, this "downward" mobility implies both switching into and away from routine tasks. Notice this is true whenever there is a movement only "downwards" or only "upwards". It is not the case if there is polarisation, or "concentration" of tasks in the middle. This is studied later.

Now, continuing our example, what is the final effect of μ_m on L_r? This depends on a complicated comparison of relative abilities of workers in the tasks they are leaving and entering into. For example, if those who are switching from routine to manual occupations are just marginally better than those who are already in manual, the expansion required on L_m to induce a given increase in T_m is larger, leading to a greater relative (and perhaps an absolute) fall in L_r. Conversely, if those who are switching from routine into manual are increasingly better than those who are already in manual, the required increase in L_m and the subsequent fall in L_r is smaller. Furthermore, this also depends on how able those who are leaving routine to manual are at routine tasks. For example, if those who switch down do not contribute much to T_r, the fall in L_r might be greater. To add further complication, the extent of switching at the bottom is greatly interlinked with the extent of switching at the top (i.e. from abstract to routine). If those switching down to routine are very good at abstract tasks, the cost of switching down is larger in terms of T_a, leading to lower switching from L_a to L_r, meaning a greater share of reallocation into L_m will come from L_r. And etc. This analysis shows how complex the evaluation of comparative statics is, even for a simple *ceteris paribus* exercise. Non-*ceteris paribus* changes are even more tricky to interpret.

The effect of σ_j on employment is less straightforward than the case of μ_j. Here, results do not depend on ϕ, except for σ_r.[24] First, notice that σ_j, just like μ_j, has a direct impact on tasks' productivity, through Λ_j (see equation 2.15). However, there is central difference between them, namely the crucial role that σ_j plays in defining comparative advantage. In effect, results in Table 2.2 indicate that the greater comparative advantage is (e.g. if σ_m falls or σ_a increases), the more employment is allocated towards the manual occupation. Conversely, the smaller comparative advantage is (e.g. if σ_m increases or σ_a falls), the more employment is allocated toward abstract tasks. The case of routine is special because a change in σ_r alters comparative advantage in opposite ways with the other tasks. As such, an increase in σ_r leads to a reduction in comparative advantage between routine and abstract, reducing employment in the latter occupations (provided $\phi \leq 0$). Similarly, an increase in σ_r leads to an increase in comparative advantage between routine and manual, fostering employment in the latter. These results are intrinsically related to the particular order of tasks along i, which is $m - r - a$. Recall that those in abstract tasks are the most skilled among all workers, whereas those at manual tasks the least skilled among workers. Greater comparative advantage induces specialisation, fostering employment in those occupations with less skilled agents, whose contribution to inputs is smaller. Conversely, lower comparative advantage induces greater concentration in the tasks where agents are relatively more skilled – abstract.

Having studied the role of ability parameters, we can reconnect the discussion with job polarisation. Table 2.2 highlights (greyed out cells) those cases that lead to JP. There are only three examples, namely (i) an increase in μ_r, with $\phi < 0$; (ii) a fall in μ_r, with $\phi > 0$; and (iii) an increase in σ_r, with $\phi > 0$. As Chapter 4 discuss, evidence seems to be very much in favour of $\phi < 0$, making the first case the most interesting one. This is, polarisation arises when workers become more skilled at the routine tasks, and substitution is relatively difficult between tasks. It is also evident that polarisation

[24] Actually, the sign of the partial derivatives depend on the **level of** i_1^* and i_2^*. This is, the derivative of the function depend on the actual equilibrium. Appendix F shows this, and discuss that in some cases the equilibrium levels are implausible, making the results presented in the table the most appealing ones.

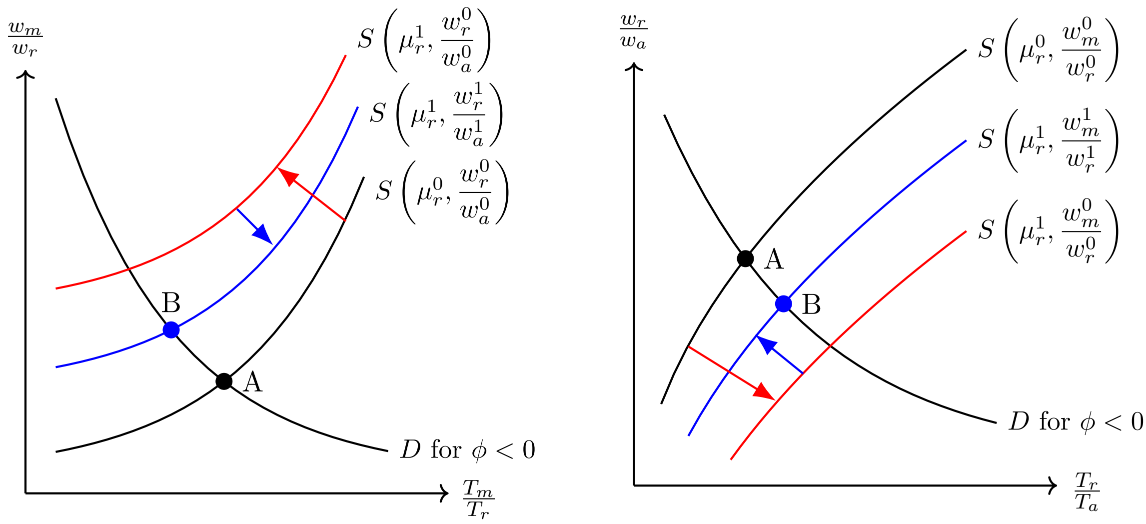

(a) Relative demand and supply for $\frac{T_m}{T_r}$

(b) Relative demand and supply for $\frac{T_r}{T_a}$

Figure 2.10: Comparative statics exercise of a relative increase in μ_r from μ_r^0 to μ_r^1, leading to job polarisation

arises when workers become **less** skilled in manual and abstract tasks, more or less proportionally.

The consequences of this polarisation process can be studied using the perspective of markets for tasks. Figure 2.10 presents the relative demand and supply for tasks, and the change in the equilibrium following an increase in mu_r from μ_r^0 to μ_r^1. Although the model is silent about the particular adjustment process towards an equilibrium, a simple sketch is shown here. Initially, every market is in equilibrium at point A. Then, there is an sudden increase in mu_r. Each relative demand for tasks is unaltered because they only depend on technology parameters (α_j and ϕ). Yet, the supply of tasks shifts in both graphs. Since workers have become more productive at routine tasks, η_r is now higher, which increases T_r for any wage rate levels. This is, the supply of $\frac{T_m}{T_r}$ falls whereas the supply of $\frac{T_r}{T_a}$ increases. The new supplies are presented in red. This change can be denominated the *ability effect*.

Now, this *ability effect* produces an excess of supply of T_r with respect to the other two tasks. This is, at the initial wage rates, firms are not demanding enough T_r. This put pressure on the relative wage rates, reducing ω_r with respect to the other. The "new" equilibrium would be those where the red lines crosses the original demand for tasks. Nonetheless, recall that supplies react to each other's prices – a "cross-price" effect.

Consequently, the new intersection points are not an equilibrium, because they assume that the relative market's wage rates are constant at the initial level. Yet, this is not longer true, as each graph indicates. In particular, in both markets w_r is expected to fall relative to other tasks, which motivate workers to switch towards manual and abstract tasks (recall Figure 2.4). This movement reduces T_r, thereby affecting the supply of tasks. For example, consider the $\frac{T_r}{T_a}$ market. Since the price in the other market is such that $\frac{w_m}{w_r}$ increases, there is switching from L_r into Lm. Therefore, there is a reduction on T_r **coming from the other market**. This must be reflect in the supply of $\frac{T_r}{T_a}$.

Following this reasoning, each market sees a fall in T_r, which pushes the new supplies in the direction of the original ones (black lines). Yet again, there is a change in relative prices (in the opposite direction), which moves the relative supplies. This process continues until there are no further changes in the supply in any market.[25] The total change due to cross-prices can be denominated the *cross-price effect*.

The final equilibrium is in point B, where the *ability effect* dominates the *cross-price effect*. Here, routine wage rate is lower relative to other tasks, but its relative output is larger, even considering the reallocation of workers away from it. This is possible because of the greater productivity of workers in routine tasks, meaning fewer yet more productive workers can still increase aggregate T_r.

Notice that this result is entirely driven by $\phi < 0$. If, for example, $\phi > 0$, the *ability effect* is the same. However, the *cross-price effect* is large enough to push the final supplies beyond the original ones. This leads to the converse story, where more employment is reallocated towards L_r, producing "job concentration". If $\phi = 0$, the two effects cancel each other.

To summarise, the simplest case of JP comes from movements in μ_r. This is consistent with the example introduced before of a widespread introduction of computers in workplaces and homes. In effect, **the change in ability should have occurred for everyone in the economy**, not only for those working in routine tasks – i.e. with computers. This is perhaps not an unreasonable postulate. Additionally, changes in

[25]The existence and stability of these equilibria have been already demonstrated in the previous subsection. Simple observation would indicate that the dynamics towards the equilibrium is that of an oscillatory process around the final supplies.

μ_m and μ_a can also be accommodated to produce job polarisation. In fact, as Chapter 4 discuss, the "bias" of polarisation depends on how the relative values of μ_j change. Regarding σ_j, its centrality in defining comparative advantage gives this parameter a special role, leading to a pattern of employment reallocation that is not likely to be by itself a source of polarisation (unless there are many changes together, again).

2.5.2 Technological change

The model used so far does not include technological change. Here, a simple formulation is introduced. Notice that we are not interested in an aggregate, neutral technical process, because this will not alter the demand for tasks; its only consequence is an increase in the final good and consumption. In our framework, this is clearly not of interest. Instead, the analysis focuses on an exogenous technical process that affects tasks' productivity individually, – the so-called task-biased technical change (TBTC). This framework is simple and intuitive enough, whilst allowing for full flexibility regarding the bias of TBTC.

Without loss of generality, let us denote $E_{j,t}$ as the productivity level of input T_j in period t. This is, the contribution of input T_j into production is now $E_{j,t}T_{j,t}$. Then, the "new" production function is:

$$Y_t = \left[\pi_m \left(E_m T_m\right)^{\phi} + \pi_r \left(E_r T_r\right)^{\phi} + \pi_a \left(E_a T_a\right)^{\phi}\right]^{\frac{1}{\phi}} \tag{2.30}$$

where π_j corresponds to the input's share in output. It follows from our assumption of constant returns to scale that $\sum \pi_j = 1$. Importantly, this function **has not changed our original model in any way**. In fact, they are identical, with $\alpha_j = \pi_j E_j^{\phi}$.

Given the new specification, it holds that:

$$\frac{\alpha_j}{\alpha_{-j}} = \left(\frac{\pi_j}{\pi_{-j}}\right)\left(\frac{E_j}{E_{-j}}\right)^{\phi} \tag{2.31}$$

This equation is very important. It tells us that the effect of TBTC on relative α_js and therefore on relative demands (e.g. equation 2.27) varies with ϕ.[26] More precisely,

[26] Input shares (π_j) are always positive (given $\phi < 1$), having no consequence on the analysis that follows.

Table 2.3: Sign of $\frac{\partial L_j^*}{\partial \alpha_j}$

Parameters	L_m^*	L_r^*	L_a^*
α_m	+	?	-
α_r	-	+	-
α_a	-	?	+

equation (2.31) implies that a relative increase in productivity of T_j with respect to T_{-j} increases (decreases) $\frac{\alpha_j}{\alpha_{-j}}$ if ϕ is positive (negative). Additionally, we know that an increase (decrease) in $\frac{\alpha_j}{\alpha_{-j}}$ raises (lowers) the demand for $\frac{T_j}{T_{-j}}$. For instance, a TBTC that improves the productivity of T_r in production raises its relative demand when tasks are relatively easy to substitute in production ($\phi > 0$); and vice-versa. Again, in the pivotal case of $\phi = 0$ (i.e. a Cobb-Douglas production function), relative demands for tasks do not change. In effect, in this scenario, any biased technological progress is actually equivalent to aggregate technical change, leading to no change in the input mix used in production.[27]

Having characterised the effect of TBTC on the model's parameters, we can proceed to study the effect of technology on employment shares. This role depends on: (i) how α_j affects i_1^* and i_2^*; and (ii) how E_j affects α_j. The latter was already described. The former is obtained from the Implicit Theorem Function, and is presented in Table 2.3 (see Appendix F for more details on the derivations). In particular, it shows the sign of the partial derivatives, equivalent to the exercise carried out for ability. The diagonal indicates that an increase in α_j expands employment in task j. Similarly to the case of ability, whenever workers' switching is "down" or "up", the effect of technical change on L_r is ambiguous. Thus, the only "pure" case that leads to JP is that of a **fall** in α_r. Naturally, combination of changes that leads to a relative fall in α_r might also conduce to JP. This is evaluated in more detail in Chapter 4.

Combining the results in Table 2.3 with those from equation 2.31, we can deduce that:

- If $\phi < 0$: job polarisation arises from a technical progress that enhances the productivity of routine tasks relative to other tasks (routine task-biased technical change, RTBTC).

[27] For example, biased technical change as in $Y = (AL)^{\alpha} K^{1-\alpha}$ is identical to unbiased technical change as in $Y = B\left(L^{\alpha} K^{1-\alpha}\right)$, with $B = A^{\alpha}$.

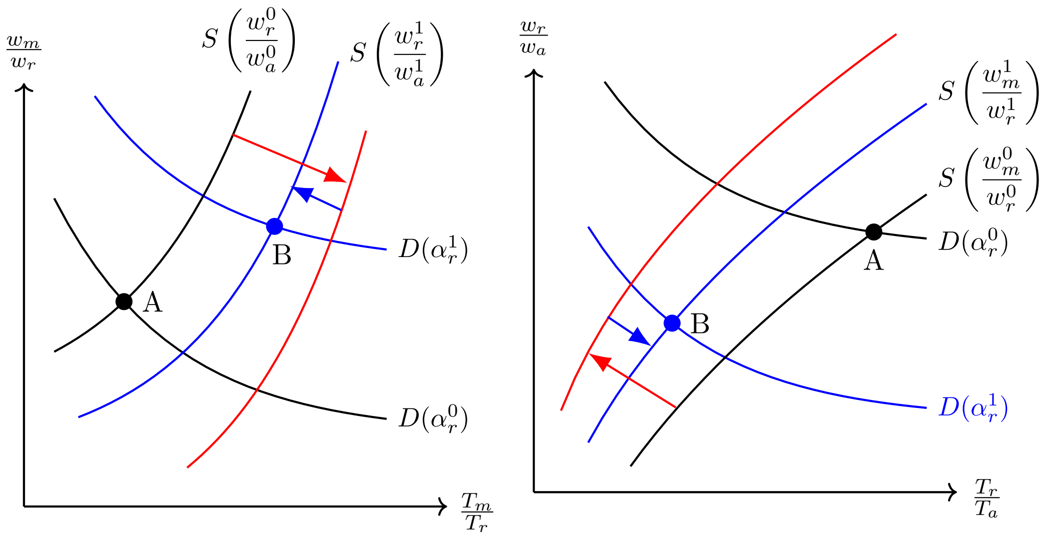

(a) Relative demand and supply for $\frac{T_m}{T_r}$

(b) Relative demand and supply for $\frac{T_r}{T_a}$

Figure 2.11: Comparative statics exercise of a relative fall in α_r from α_r^0 to α_r^1, leading to job polarisation

- If $\phi > 0$: job polarisation arises from a technical progress that enhances both manual and abstract tasks relative to routine tasks.

As Chapter 4 discusses, $\phi < 0$ is empirically the most compelling scenario.

Let us now evaluate the case of RTBTC using the markets for relative tasks. Given $\phi < 0$, this translates into a fall in α_r – say from α_r^0 to α_r^1. Figure 2.11 presents the state of these markets before and after the change. The initial equilibrium is in point A. As said, RTBTC leads to an immediate change in the relative demands, shown as blue lines. Importantly, without further change in technological parameters, these are the demands that prevail in the final equilibrium. This change in demands can be named the *productivity effect*.[28] Regarding supplies of relative tasks, these do not suffer an "immediate" or direct shift, because they depend on ability parameters and the other market's relative wages, but not on final good's technology.

Now, "after" the demands shift, the original equilibrium no longer holds. There is an excess demand for both T_m and T_a, equivalent to an excess supply for T_r. These disequilibria induces the relative wage rate for routine tasks to fall. This change in

[28]This is quite different from the *ability effect*, which refers to the change in workers' ability for performing a given task. The *productivity effect* refers to how productive inputs are on the production of the final good. In other words, the former affects the supply of tasks whereas the latter affects the demand for tasks. This is evident from the graphical analysis.

“cross-prices” is subsequently reflected on each market, shifting the supply curve. The direction of change is identical to that related to changes in ability parameters. Namely, the reallocation of employment in the other market changes relative supply of tasks in that other market, thereby affecting relative tasks in the own market. In the example here, for a given level of $\frac{\omega_m}{\omega_r}$, there is an increase in $\frac{T_m}{T_r}$, shifting the supply in Panel (a) to the right. From a similar argument, the supply in Panel (b) shifts to the left. These are shown in red.

The dynamics that follow mimic the ones in Figure 2.10. The *cross-price effect* realises fully at the final supply curves, in blue. The final equilibrium is in B. Here, the routine wage rate **and** the output of routine tasks falls relative to manual and abstract tasks. The intuition of this was already noted. The nature of TBTC assumed here (RTBTC), together with the relatively low elasticity of substitution between tasks, means production of Y expands the most by moving workers towards the least productive inputs, T_m and T_a. Since the population size is fixed, the only way this comes about is through an increase in L_m and L_a (i.e. job polarisation).

There is an interesting contrast with respect to changes in ability parameters. There the **ability effect** and *cross-price effect* move over the same dimension. This is, they both move relative prices and relative tasks in opposite directions (this is, **along** the demand). In turn, the only possible result is that of an increase (fall) in the relative value of ω_j and a fall (increase) in the relative input T_j. In contrast, here the *productivity effect* and the *cross-price effect* fully define the complete movement space (as changes occur along the demand and supply). In turn, **any** final result in terms of relative wages and inputs is possible.

The above analysis focused on a “symmetric” form of RTBTC, because it only changed productivity of T_r. A more general form of technical change that, for example, lowers the demand for routine tasks with respect to manual and abstract tasks, **and** increases the demand of abstract tasks with respect to manual and routine tasks (e.g. if α_r falls and α_a increases), might also lead to polarisation, albeit not always. In effect, if the bias with respect to abstract tasks is relatively low, job polarisation can still occur. However, is the bias is big enough, it could be optimal for firms to reduce the output of both T_m and T_r, in favour of T_a. The conclusion from this thought experiment is that there is no other generalisation of JP, besides that of "pure" RTBTC. Any other case can only be resolved for a particular parametrisation.

At this point it might be illustrative to reconnect with the main argument in Autor and Dorn (2013). They assume an exogenous fall in the price of capital, which induces firms to increase its use in production. This enhances the productivity of routine workers. As such, their story is the same that the one here, of an exogenous increase in E_r. Their model however has two goods, both which are valued for consumers. As such, there is a trade-off between the elasticity of substitution in production and consumption. In their case, RTBTC might not lead to polarisation if substitution in consumption is comparatively high; consumers would prefer to consume more of the good produced with capital, thereby inducing greater employment into routine tasks. This is clearly interesting. Additionally, their model allows for RTBTC to induce polarisation even if $\sigma_Y > 1$, which here is not possible. However, the evidence contradicts these two facts. Namely, polarisation exists, and $\sigma_Y < 1$. As such, the story here, albeit more simple, is enough to capture the important developments observed in advanced economies, based on more realistic assumptions.[29]

2.5.3 Difference across countries

Figure 2.12 reproduces that of Figure 0.4 in the introductory chapter. This shows that advanced economies present diverse patterns of employment reallocation over the same period. In particular, there is difference in terms of (i) degree of "hollowing out" (fall in routine employment), and (ii) the bias in the reallocation, e.g. mainly upward, downwards, or evenly distributed. It is clear from the previous analysis that such diverse phenomenon can be entirely captured by this model, by simply shifting the parameters accordingly. In fact, an approach like this leads to an undetermined system, as the number of parameters to choose is much greater than the number of employment shares to fit.

[29] Furthermore, next chapter shows that the distributions of ability are very far apart from those in Autor and Dorn (2013).

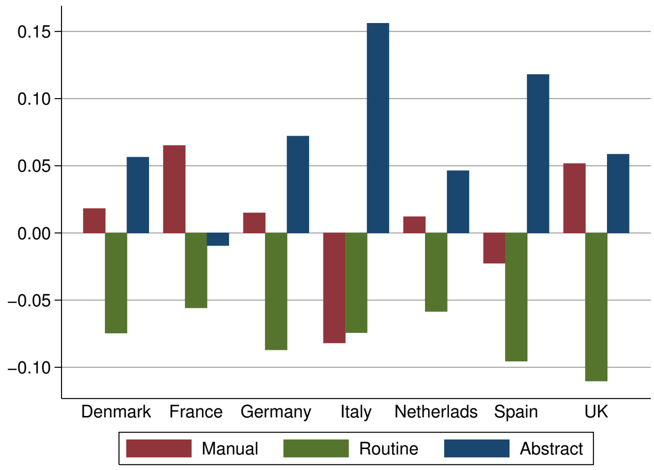

Figure 2.12: Change in employment shares, 1992-2008, selected EU countries

Still, there is some room for a relatively general analysis. For example, Figure 2.12 presented the case of United Kingdom, which is an example of strong hollowing out of employment, with fairly equivalent reallocation up and down. This could mean the UK might have faced a TBTC that did not affect relative productivity between manual and abstract inputs in production but did affect routine one, perhaps together with some changes in the distribution of ability that did foster manual or abstract workers' skills relatively even. Chapter 4 studies the case of the UK in more detail, finding a strong role of technology on polarisation, and little of abilities. This result is permitted by the estimation of the ability distributions for the UK – in Chapter 3, which reduces considerably the indeterminacy of the system mentioned above.[30]

This might contrast with the case of France, which shows around half of the hollowing out of the UK, and a strong reallocation towards manual occupations. In our model, this could be due to a particular form of TBTC that greatly affect manual and routine tasks relative to abstract. For instance, if $\phi < 0$, there might have been an increase in productivity of routine tasks and abstract tasks relative to manual, which then fosters the reallocation of employment towards the least productive tasks. Alternatively, or perhaps jointly, workers might have become more able at both routine and abstract tasks. Chapter 4 also studies the case of France, by testing the hypothesis that the technological process affecting the UK might have been the same that the one affecting

[30]It is shown that there is only one free parameter, ϕ, which is calibrated from existing empirical literature.

France (perhaps due to a global technological shock). The exercise finds that even if the shock is global, the consequences in terms of productivity of T_j must differ, as changes in ability distribution alone cannot explain such stark contrast between UK and France.

The case of Germany is the opposite of France. Here, reallocation of employment has been mainly towards abstract occupations. Again, there is indeterminacy, as this could come about through a combination of multiple changes in workers' ability and inputs' productivity.

The overall conclusion is that both ability and general productivity of tasks in production could lead to drastic changes in employment. The exact magnitude and bias of employment shares shift is impossible to generalise, specially because the model does not offer closed-form solutions. Still, it is clear the framework is flexible enough to reproduce any observed pattern. Chapter 4 offers more insights on the fitting of the model to real data, thereby demonstrating its applicability to the study of job polarisation.

2.6 Conclusion

This chapter presented a task-based model of sorting and comparative advantage, where workers have heterogeneous ability for different tasks. It enables a very precise characterisation of the role these distributions have in employment allocation in a given period, and over time. The latter is crucial to understand phenomena like job polarisation. As such, this chapter contributes to the literature by proposing this new mechanism, which so far the polarisation literature has overlooked. It highlights the distinct role both comparative and absolute advantage play on workers' reallocation and job polarisation.

The model developed here is central in the next two chapters. Given the predominance of the ability distributions for the whole analysis, the next chapter set out the goal to estimate these distributions for the UK. This provides a calibration to this important component of the model, which is used in Chapter 4 to identify the nature of the technological process affecting the UK economy.

There are certainly areas in which the model could be expanded. Perhaps one very profitable extension is to add an external sector. For example, the introduction of a second country playing the role of China or East Asia, as a place to which part of the production of routine tasks or inputs have moved. There is already some literature on this – related to off-shoring. It would be possible to connect them. The benefit of such additions introduces some essential mechanisms from trade theory. In fact, not only there is comparative advantage between workers within an economy, but also between workers across countries. This might add important insights about the role ability and TBTC plays on polarisation.

Other extensions possible, which seem to be empirically relevant are mobility costs, perhaps related to learning and occupational-specific experience. The switching incentives would change under these scenarios. Also, the introduction of a state variable like experience would add greater dynamic to the model. Another benefit of adding switching costs is the evaluation of misallocation of resources, which is known to occur in both developing and advanced economies alike (e.g. Asker, Collard-Wexler and Loecker, 2014). It might also be interesting to detach tasks from occupations. The motivation for this is again empirical, as many occupations seems to be a combination of tasks rather than be mainly compromised of a single task (e.g. Bisello, 2013). In any case, all these risk over-complicating the model, and it is not clear whether the new insights might balance the added obscurity into the model's mechanisms. All these are areas that I expect to analyse in my future research.

3

Estimation of the skills distribution for the UK

3.1 Introduction

Chapter 2 showed the important role the distributions of ability play in understanding employment reallocations due to technological forces. Thus, the empirical characterisation of these distributions in a single period and over time becomes crucial if one is to understand observed labour market developments in a given country. This chapter estimates these distributions for the UK workforce. here, consistent with the model of Chapter 2, workers are assumed to posses an ability for manual, routine, and abstract tasks. The ultimate objective is then to characterise workers' skill set for **all** tasks.

Ability is by definition unobserved (by the econometrician). Yet, for our purposes, it is useful to differentiate between ability in occupations where workers are observed (i.e. employed), and ability in occupations where they are not observed. This differentiation is central to this chapter because the method used to estimate ability differs between them. In effect, the latter is based on the results from the former. Thus, this chapter is structured in two parts. The first part (the longest) focuses on estimating ability in "observed" occupations. The second part imputes ability in "unobserved" occupations.

Ability in "observed" occupations is estimated from wages, using a panel-data Mincer regression. This regression is derived from a simple selection model, where workers sort into occupations based on comparative advantage. Since wages are expected to reflect workers' ability, by controlling for all observed factors known to determine earnings (like education, experience, gender, etc), residual variation in wages that is specific to the individual – i.e. unobserved heterogeneity – proxies for ability. This econometric model is estimated using data from the British Household Panel Survey, for the period 1991-2008. The estimation is carried out using four methods, namely Random Effects (RE), Fixed Effects (FE), Hausman and Taylor (HT), and two stages least squares fixed effects (FE-2SLS). FE and HT allow to control for potential correlation between education and ability, whereas FE-2SLS controls also for endogeneity of education with respect to the idiosyncratic error – for example, due to simultaneity bias. The comparison between methods provides interesting insights about underlying endogeneity.

The econometric model set out here has one particularity. Since ability varies across occupations, using a standard panel-data model is not enough, because the latter has an unobserved component which is specific to the individual i – c_i, independent of occupation. The solution is to redefine the clustering of the model. Instead of worker i, a new "occupation spell" cluster is defined, where every worker that switches occupations is considered as a new observation unit. This new model has an unobserved component that varies across individuals **and** tasks. Key to the correct identification of ability in this framework is the assumption that no other individual-level unobservable varies across occupations.

Two important steps are taken before carrying out the estimations. First, occupations are mapped into tasks. In effect, survey data includes information about workers' occupations, based on some common classification standards. Yet, the model approaches ability from the tasks framework. As such, it is necessary to link these two. The mapping is based on the existing literature, albeit adapted to the particular dataset. Since there are at least two occupational codes standards available in the sample for every worker, a more refined mapping is conducted by lowering incoherences between these standards.

The second important steps refers to sample selection evaluation. Inference from the sample to the wider population requires selection to be only based on observable variables.

In this case, there is evidence that selection does depend on unobservables. In particular, there is evidence of attrition, although female labour self-selection seems to be low. These results (presented in Appendix G) play against our wider inference attempt.

The model's estimation reveals that workers' ability in observed occupations is positive skewed (e.g. like log-normal distribution), for each task. There is also evidence that average ability has been falling over the period, driven by lower ability of those who enter the sample, higher of those who exit, and also due to the nature of switching. Still, the wage rates paid to both manual and abstract tasks has increased compared to routine tasks, in accordance with the model in the next chapter.

The analysis just described produces estimates of ability for just half of the workers' skill set. The second part of this chapter consists of producing an estimation of ability in occupations where workers are not employed – the remaining half. Because there is no information about workers' wages in the latter, an observational approach like a regression is not helpful. In order to impute this unobserved ability, a theory-based approach is needed. Here, a simple labour market model is assumed. Given a competitive labour market and perfect information, workers self-select into occupations, choosing that which maximises their current income, conditional on their skill-set. Thus, under these conditions, workers' observed occupational decision reveal information about their remaining unobserved skills. A relatively simple empirical analysis is able to provide good support to some of this model's assumptions. More precisely, there is evidence that (i) after considering switching costs and information problems, most of workers choose the occupation that maximises their wage, and (ii) there is a high degree of absolute advantage between workers' ability. This is, workers that are good at one task, are generally good at other tasks too, and vice-versa.

The central insight arising from the model that is helpful in identifying unobserved ability is that workers are not observed in a given occupation because their ability for them is not high enough to motivate them to switch. In practice, this imposes a "bound" to all unobserved abilities. Workers ability could be anywhere within zero and this bound. A first step forward would simply assume missing ability to be equal to the boundary. Another would assume ability to be halfway between zero and the boundary.

In any case, these two "extreme" examples show that the complete ability distributions are positive skewed – as it is for any in between these extremes. This is taken to be the conclusion from the exercise.

There is an interesting remark that arises from this exercise. The approach used here to estimate ability is *agnostic* in the sense that there is no ex-ante structure imposed to the nature of ability distributions. This is, distributions could differ in any way across tasks. This contrasts with approaches based on Lemieux (1998), which assume the nature of the distributions to be identical across the dimension of interest (occupation, sector, etc), albeit allowing for moments to differ between them.[1] As the results here indicate, the unrestricted approach favoured here cannot take us very far in the goal of characterising true, underlying distributions, even based on a stylised model. It is shown that this is an intrinsic feature of the data generating process, and no much further can be said, unless more structure is imposed on unknown distributions. Because an agnostic approach was preferred here, the conclusions might look rather limited. Yet, in our view, they are still interesting. Less agnostic approaches – like that in Lemieux (1998) and derivatives – can certainly be complementary to those here, and future work is expected to be conducted in this area.

Several contributions to the literature arise from this chapter. First, as the literature review in the next section highlights, existing attempts to estimate ability distributions from wage regressions only focus on occupations or sectors where workers are observed. Additionally, some of them impose strict structure on the nature of ability distributions. This paper fills this gap by providing the first attempt to estimate the **complete** ability distributions of a workforce using unobserved heterogeneity from wage regressions. These distributions also include ability in occupations where workers are not observed. The method used here is also more agnostic, as it does not rely on ex-ante distributional assumptions. There are of course caveats and strong assumptions involved in the process, and room for improvement does exist. Still, this chapter provides an important step forward in that direction.

Second, the empirical results seem consistent with two labour market phenomena, namely workers' optimisation and absolute advantage across workers' ability, upon which many

[1] For example, ability in one sector distributes $\sim \mathrm{D}(\theta_1, \sigma^2_{\theta_1})$, whereas in another sector distributes $\sim \mathrm{D}(\theta_2, \sigma^2_{\theta_2})$, with D being the same.

sorting models are based (e.g. Lemieux, 1998 and related literature). Although the fit is not perfect (e.g. switching costs and information problems might exist), these results seems to be convincing enough to make them an interesting empirical contribution.

Finally, a seemingly novel test is implemented to test the central tenet of the methodology, namely that unobserved heterogeneity varies across occupations. This Hausman-based test seems to be useful not only for this particular task-based model but for any where unobserved heterogeneity varies across a second dimension (like occupation, union, sector, etc). The referenced literature also provides evidence that heterogeneity varies across a second dimension. Here a formal test is conducted to evaluate this hypothesis, finding strong support in favour of multidimensional heterogeneity. This has important consequences for the Mincer literature, which usually relies on unidimensional heterogeneity. It is shown that the ubiquitous Fixed Effects methodology is inconsistent if heterogeneity is multidimensional.

An important note on terminology

The literature on wage regressions usually refer to ability implicit in the unobserved heterogeneity component as “unobserved” ability (to the econometrician), in contrast with “observed” ability like education, experience, etc. The notation here is different. True, the interest is in characterising the “unobserved” ability distribution, but the terms observed and unobserved here refer to whether the worker is observed or not **in a given occupation**. Recall the objective of this work is to characterise the complete skill set of workers. Yet, the majority of them are employed only in one occupation. Since the wage regression approach can only recover ability from occupations in which workers are observed (see graphs in next section), there is a need to distinguish between occupations where workers are employed versus those at which they are not employed. Hence the notation “observed” ability (recovered from observed occupations) versus “unobserved” ability (recovered from unobserved occupations).

3.1.1 Road map

This Chapter continues as follows. Section 3.2 provides a brief context on the methods used to estimate ability. Section 3.3 presents the econometric model to be estimated. The next section describes the four estimation methods applied to the econometric model. Section 3.5 briefly describes the dataset used, whereas Section 3.6 explains the mapping between occupations and tasks, essential for the analysis. Section 3.7 goes through the main model estimation, presenting the results relating to observed ability. Section 3.8 takes the inputs from the previous section to estimates unobserved ability with the help of the model developed in Chapter 2. To justify this method, the model's assumptions are evaluated. Finally, Section 3.9 concludes.

3.2 Literature review

Ability from unobserved heterogeneity on wage regressions

One method to produce estimates of ability (or productivity) comes from wage regressions. This usually comes from a setting where individuals' ability varies across certain dimension (e.g. public/private sector). If individuals sort into groups based on ability, this selection needs to be accounted for when estimating causal effects. The latter is usually the objective of this literature, rather than the characterisation of ability. Thus, in these studies, the subsequent analysis of derived ability distributions is either missing or incomplete. This chapter – which also uses the wage regressions approach – seems to be the first directed attempt to provide a complete characterisation of the ability distribution of a sample or population.

The seminal paper in this literature is Lemieux (1998), who estimates the role of unions on wages, using panel data for the US. From our perspective, what is interesting about this paper is the characterisation of ability, used in subsequent papers. The economy has two sectors, the union (U) and non-union (N) sector, where workers sort based on comparative advantage. Ability of workers is given by $\theta_i^N = \theta_i + \xi_i$ and $\theta_i^U = \psi\theta_i + \xi_i$ respectively, where θ_i is a "deep" ability parameter and ξ_i is an error term. Notice that this specification assumes the ability distribution is the same in both sectors, given by the distribution of θ_i. The moments of the distribution might differ though, provided $\psi \neq 1$. This constrained approach contrasts with the agnostic method used here, where distributions might differ across tasks. Lemieux then goes to estimate such a model, using a GMM approach, but the analysis does not discuss the resulting estimation of θ_i, even though its identification was possible.

Lemieux's framework is used by Suri (2011), who is interested in identifying the benefits from advanced technology adoption among farmers in Kenya. To account for selection, she models the sorting of farmers into production techniques, sorting which is based on comparative advantage over farmers' productivity for different technologies (the equivalent to ability in our setting). In this case, there are two technologies, non-hybrid (N) and hybrid (H). The productivity of farmers in both is $\theta_i^N = \theta_i + \tau_i$ and $\theta_i^H = (\phi + 1)\theta_i + \tau_i$. Again, this framework assumes that the underlying productivity distributions are identical in their form. After conducting the main model's estimation, the author set out to describe the distributions of productivity arising from the estimations.[2] She differentiates between farmers using hybrid (H) and non-hybrid (N) methods during the two sample years, and between those who are only one year in the sample (leavers L and joiners J), regardless of their farming technique. The resulting distributions (*pdf*) are reproduced in Figure 3.1, equivalent to Figure 5C in the original paper. Importantly, these include **only** the productivity of the method farmers are observed using. For example, if a farmer is using hybrid methods in both years, he or she is only represented in the H distribution. As such, these are not the complete distributions, as this paper aims to characterise. That is the reason why they are not equivalent (as per assumption). Finally, the paper does not goes into a parametric characterisation of these distribution, nor does it attempt to identify the population distributions.

Imbert (2013) is another example of careful selection consideration, following Lemieux (1998) framework. In this case, the subgroups are public and private sectors, in a model applied to Vietnamese data. The definition of θ_i here is that of an index of comparative advantage composed of "all of the productive characteristics that influence a worker's

[2]The estimation of θ_i comes from using a Chamberlain's decomposition (based on Chamberlain, 1982), where θ_i is defined as a linear projection of the histories of farming techniques used – which are observed. See section 4.5.1 in paper for more details.

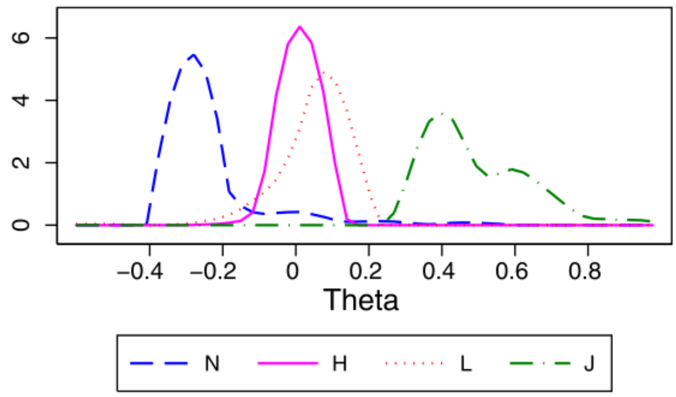

Figure 3.1: "Observed" productivity of farmers, from Suri (2011), reproduced with permission of the Econometric Society

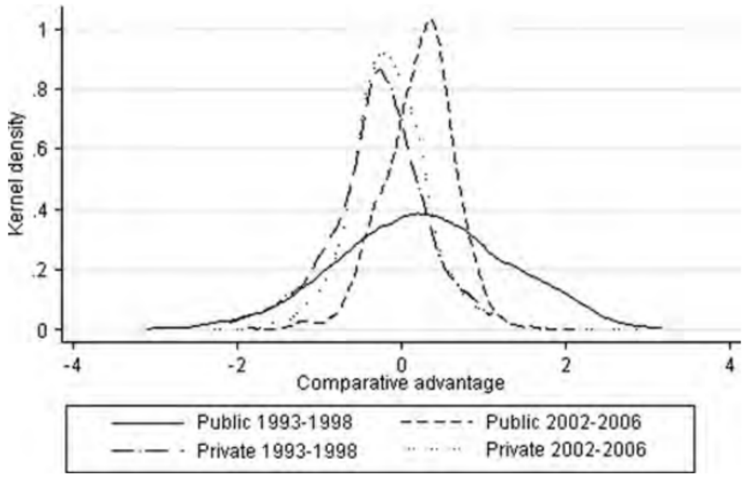

Figure 3.2: "Observed" productivity of workers, from Imbert (2013), reproduced with permission of the Oxford University Press

pay in one sector relative to the other" (p.61). Private and public productivity of worker i are θ_i and $\alpha\theta_i$. The model is then also assuming equal distribution (with different moments, if $\alpha \neq 1$). The author proceeds to estimate individual heterogeneity after carrying out the main model's estimations. These are shown in Figure 3.2, equivalent to Figure 3 in the original paper. The shape of these distributions is not clear – between low skewness to relatively clear negative skewness. Again, these correspond to ability of workers in observed sectors, reason why they are not of the same nature, as the underlying distributions are. Just like in Suri (2011), the analysis stops here; no attempt is made to recover population distributions, or to characterise their parameters.

In summary, these papers set out a framework from where a common distribution of ability can be obtained, a distribution which varies across a single dimension like union status, sector, or technology. Yet, because of other objectives, this framework has not been used for the complete characterisation of the population's underlying ability distributions.

To the best of my knowledge, there seems to be no attempt to apply Lemieux's framework to (three) tasks, even though this is clearly possible – with the aforementioned limitation of equal distribution. This method is explored later in this chapter. There is however one paper – Cortes (2016) – that proposes a sorting model with three tasks, which then can be used to characterise ability distributions. His approach is different from that of Lemieux. Instead of introducing tasks as a variable in the model – like authors above use dummies for union/non-union, hybrid/non-hybrid, or public/private, Cortes introduces another dimension to the econometric model. In effect, the papers reviewed above use individual and period (i and t) as panel dimensions. Cortes adds a third – occupation j. The benefit of this structure is that unobserved heterogeneity – normally individual specific – can now vary across i and j. This is, unobserved heterogeneity is occupation-specific. The implementation of this method requires to redefine the clustering of the model, by treating workers who switch occupations as different units. The benefit of this approach is the agnostic nature of unobserved ability and the underlying distributions.[3] Cortes sets out to estimate the model for the US. The estimation of ability are presented in Figure 3.3, which reproduces Figure A.11 in the paper's web appendix. Like before, these only represent ability from "observed" occupations, and not the complete ability set. Yet, this is the closest existing paper to the analysis presented in this chapter.

The current gap in estimating population-wide distributions is even more evident from the theory side, at least with respect to tasks. Autor and Dorn (2013) provides a good example of this. They set out an interesting model of job polarisation. Yet, the assume that ability in manual and abstract tasks is homogeneous (i.e. the distribution is degenerate), whereas ability in routine tasks is exponential. In light of the evidence shown so far – and this chapter's results, those assumptions are quite unrealistic. They

[3]There is certainly some correspondence between Lemieux and Cortes approaches. In a sense, any variable characterising a discrete choice (like union, sector, industry, etc) could be thought as an extra dimension of the panel, and vice-versa (for example, as in the multilevel, or hierarchical modelling literature, e.g. Snijders and Bosker, 2011). Yet, a more formal claim that they are approximately equivalent is not to be implied here, but it is surely an interesting point for future research.

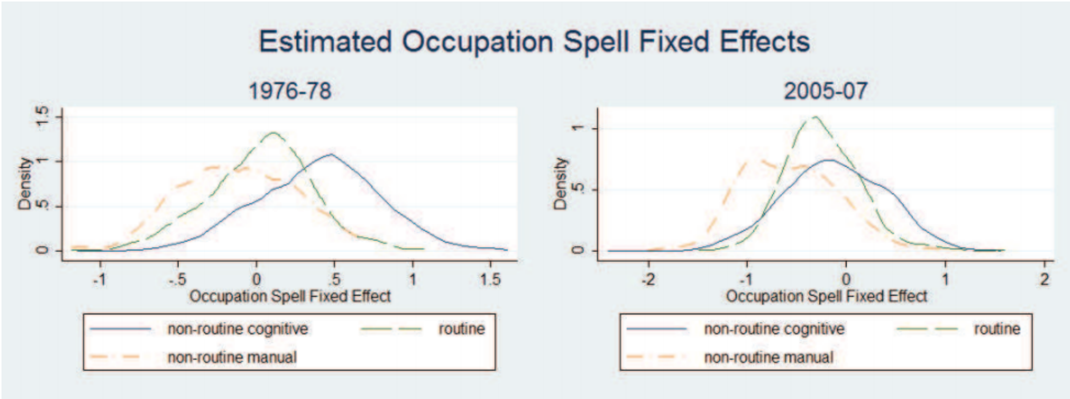

Figure 3.3: "Observed" productivity of workers, from Cortes (2016), reproduced with permission of the University of Chicago Press

also contradict absolute advantage, which this chapter's results support. All of this demonstrates the relevance of this chapter's goal for the existing literature.

Ability from tests

There is of course a large literature estimating ability from tests, particularly in fields like education and psychology. There is also an emerging literature in economics. A review of this literature is beyond the scope of this chapter. See Heckman and Kautz (2013) for a comprehensive survey regarding economics. Still, an example of what this approach aims to achieve is insightful, and allows us to mention some of its limitations.

Heckman, Humphries and Veramendi (2016) estimates the effect of education on wages in a comprehensive, dynamic discrete choice model where previous education decisions open up new educational stages. They account for selection by modelling the role of individuals' ability endowment in their educational decisions. Ability also affects outcomes, as usual. They assume agents have two types of endowment, cognitive (θ^C) and non-cognitive or socio-emotional (θ^{SE}) ability. They obtain estimates for these endowments based on the data from two tests, an arithmetic reasoning test (ASVAB) and a language arts grade measure. Tests' results are assumed to depend on these endowments and other covariates. The key identification assumption for this two equation system is that θ^C only affects the ASVAB test.[4] The resulting ability

[4] For more details, see section A.4 of web appendix.

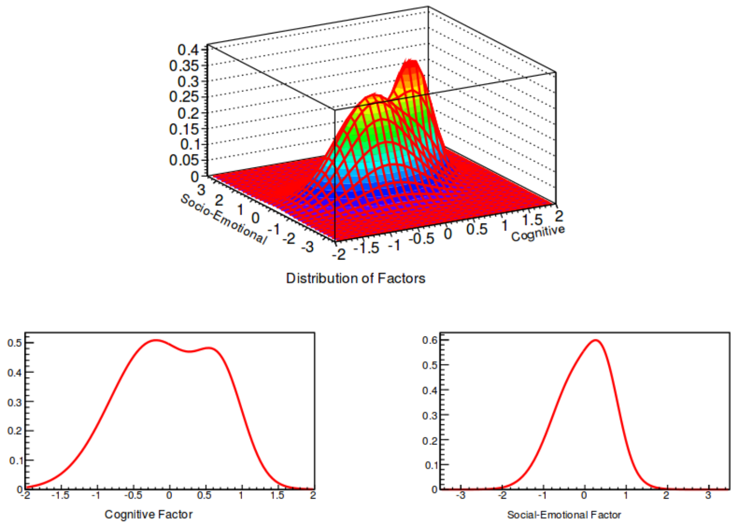

Figure 3.4: Ability from tests among sample of US population, from Heckman, Humphries and Veramendi (2016), reproduced with authors' permission

endowments are shown in Figure 3.4, equivalent to Figure A.4 in the web appendix.[5] These show a pattern that might be similar to some of the results already shown, namely an asymmetric distribution with negative skewness. The positive correlation between these tasks is of interest. Notice that these estimates correspond to the whole sample. There is no distinction between unobserved and observed dimensions. This contrasts with the wage equation method, where estimates correspond only to "observed" or selected samples.

There are however several problems with estimating ability from tests. First, the definition of ability is narrower, albeit more precise. In the example above, cognitive ability refers to particular skills captured in a mathematical test. Alternatively, ability might come from IQ estimations (e.g. Nisbett et al., 2012). This contrasts with the literature of ability from wage equations, where the definition is much more practical, and it usually refers to "all unobserved factors affecting productivity". The benefit of the latter approach ensures a direct relation between "ability" and wage. Tests-based inference however does not ensure a direct relation between ability and wage. Recall that what is relevant for sorting is comparative advantage on factors that are relevant for productivity and therefore for wages. It is unclear, for instance, how socio-emotional ability affects productivity and comparative advantage across occupations.

[5]Another example of a similar analysis can be found in Carneiro, Hansen and Heckman (2003), Figure 2.

The drawback of the wage approach is the often unverifiable assumption that no other time-invariant unobserved factor but ability is captured by unobserved heterogeneity.

There is a second problem with estimating ability from tests, that in my view limits considerably the use of this method. This is the intrinsic endogeneity of tests' grading scales. As Atkinson (1983, p.199) brilliantly puts it:

> *Suppose that the tests consists of ten questions, each involving the addition of two numbers. Most people are likely to score close to 100 per cent on such a test. On the other hand if the ten questions involved solving differential equations, most of the population would score zero. By varying the ratio of easy and difficult questions, we can get almost any distribution that we like. The fact that most actual IQ tests lead to a distribution of scores that follows the normal distribution does not necessarily tell us anything, therefore, about the distribution of abilities; it might simply reflect the way in which the tests have been constructed.*

This critique also limits tests' comparability. Results from different tests might lead to conclude that ability varies across concepts, or population, or whatever makes them different, when they can actually reflect idiosyncratic difficulty and/or grading systems. Even if full information on each test's details are available, adjusting for them is likely to be a daunting task. None of these critiques apply to the wage method. Certainly, the distributions might differ based on the definition of variables, but if sorting is indeed based on comparative advantage, the same result should arise in that respect.

3.3 Econometric model

The econometric model presented here is an extension of the theoretical model presented in the previous chapter. Notably, factors other than the wage rate and ability are allowed to affect wages, including an idiosyncratic error. Also, a time dimension is introduced into the model.

Wage determination

Assume the wage of worker i in occupation j in period t is given by the following equation:[6]

$$\ln w_{ijt} = \ln \omega_{jt} + x_{it}\alpha + z_i\beta + \ln \eta_{ij} + \epsilon_{it} \tag{3.1}$$

where w_{ij} is the wage a worker i would receive if employed in occupation j – in real monetary terms, ω_j is the "wage per ability unit" in occupation j – also denoted wage rate, equal for all workers, and η_{ij} is worker's ability in occupation j. x_{it} and z_i are other worker characteristics (discussed at length in Section 3.5), time-varying and time invariant respectively. ϵ_{it} is an idiosyncratic error, with $E(\epsilon_{it}) = 0$ and $V(\epsilon_{it}) = \sigma_i^2$. Notice that wage rates might vary over time, whereas ability is fixed. Importantly, ability and experience are separate factors. The latter can vary over time, and it is assumed to be general to each occupation. As such, experience (and further powers of experience) belongs to x_{it}.

This model allows a rich variety of mechanisms determining wages in an economy. Product and labour demand and supply factors affecting wages (e.g. technological change, or migration flows) are captured in ω_{jt}. The supply of labour also depends on the aggregate distribution of workers' characteristics like ability, education, etc. Wages are a direct function of ability, meaning pay is directly related to workers' marginal productivity at work. Other workers' characteristics affect wages via mechanisms like productivity (e.g. education and experience), discrimination (e.g. gender), geographical differences (e.g. region), bargaining power (e.g. union membership), etc. Naturally, a general equilibrium setting would model all elements already mentioned. The labour market equilibrium would be achieved via movements in efficiency wages ω_j. For the purposes of identifying ability distributions however, it is enough to model the wage rate as an exogenous process.

This model is sometimes known in the literature as Mincer equation, in honour of Mincer (1974), who first specified this functional form (although in a cross-section

[6] Observationally, to identified a worker is enough to use subscripts i and t. Thus, we can use w_{it} instead of w_{ijt}. This assumes a worker cannot hold two occupations in one observation in the sample. While some workers do report having more than one job, these are eliminated from the sample because the dependent variable – defined as hourly wage – is cumbersome to calculate for them.

setting). In this particular context, equation (3.1) is a panel data model, with two unobserved components, namely the log of wage rate and the log of ability. Unlike standard panel data models, here both unobserved components vary over occupations. This feature introduces considerable estimation issues. In order to obtain a standard panel specification, two modifications are then needed. The first one takes advantage of the fact that the wage rate does not vary over individuals. Since there are only three occupation groups, and T periods, there are $3 \times T$ wage rates to identify, a relatively low number. Then, these can be identified using a set of occupation-year interaction dummies. More precisely, define D_{ijt} as an indicator variable equal to one if individual i is in occupation j in period t. The modified model is:

$$\ln(w_{ijt}) = \sum_{j\times T} D_{ijt}\ln(\omega_{jt}) + x_{it}\alpha + z_i\beta + \ln(\eta_{ij}) + \epsilon_{it} \tag{3.2}$$

The coefficient of each of these $3 \times T$ new variables correspond to a respective $\ln(\omega_{jt})$.

Unfortunately, we cannot follow the same approach for ability, because, unlike the wage rate, the former varies over individuals.[7] Instead, we can make use of the fact that workers are observed in only one occupation in order to create a new index, $I(i,j)$, which uniquely represents a worker-occupation pair. Thus, whereas i identifies a given worker regardless of its occupation, $I(i,j)$ treats the same worker in different occupations as different units. This new panel or cluster level is known in the literature as occupation-spells (e.g. Cortes, 2016) because it treats units as part of the cluster as long as they belong to the same occupation. Notice the new clustering still has the same number of observations, N. All it does is to remap i to I.

Under this new definition, we can arrive to the **final model**:

$$\ln(w_{It}) = x_{It}\alpha + z_I\beta + \sum_{j\times T} D_{It}\ln(\omega_{jt}) + c_I + \epsilon_{It} \tag{3.3}$$

where, for notation simplicity, I have replaced $\ln(\eta_I)$ with c_I. Equation (3.3) is a linear panel model with standard unobserved effects, as c_I is unique to the individual within occupation spells. Therefore, it can be estimated using standard panel data methods, from where a prediction for each c_I can be obtained.

[7] We would need, for **every** individual, one dummies for each occupation in which he or she is observed. This, in principle possible, is highly inefficient, and might be computationally impossible.

This new model has a caveat though. The new clustering does not alter the error structure in terms of mean ($E(\epsilon_{It}) = 0$), heteroskedasticity ($E(\epsilon_{It}) = 0$), and serial correlation, but it introduces cross-sectional dependence (CSD), as both the idiosyncratic error and ability for different tasks within individuals might not be random. This is, $\epsilon_{It} \not\perp \epsilon_{I't}$ and $c_I \not\perp c_{I'}$ respectively, for same i in I and I' (i.e. switchers). CSD is problematic because consistency of common panel data estimators require independence between observations.[8] This is specially important in RE and HT models, where the composite error includes both c_I and ϵ_{It} (recall FE eliminates c_I). In our case however, the extent of CSD is very low. As shown later, only 44% of individuals in the sample are switchers, meaning CSD is not a problem for at least 56% of the sample (because $i = I$). Additionally, for switchers, the correlation problem is only between two or three I clusters (depending on whether they switch between two or all tasks), out of a large N. This contrasts sharply with cases of CSD where spatial effects (like local labour markets) or aggregate factors (e.g. house price shocks) lead to correlation between most or all of individuals with most of all the other individuals. Borrowing the terminology from Chudik, Pesaran and Tosetti (2011), our re-clustering is likely to introduce *weak* CSD, – in the sense that the variance of the error term across I goes to zero as $N \to \infty$, for any period. Although a formal proof is required for this, the intuition is that growing N lowers the aforementioned 2 or 3 I to N correlation faster than what the increase in new individuals for which correlation between 2 or 3 I exists. This is because the former affects all in the sample, whereas the latter refers to new individuals only, whose correlation is also falling because of the higher N. In turn, the expectation here is that CSD is not really a problem for our estimation. The same approach seems to be followed by Cortes (2016), who is also not concerned by CSD.

Unbiased and consistent estimation of c_I

Since all the analysis to be carried out here builds on c_I, it is central to obtain an unbiased (in short panels) and consistent (in long panels) estimation of it. As shown here, this is only possible if we can obtain an unbiased and consistent estimation of the model's parameters, namely α, β, and every wage rate $\ln(\omega_{jt})$. A description of the

[8] For more information about CSD, see Chudik, Pesaran and Tosetti (2011).

methods used to obtain these parameters is left for the following section. For now, it suffices to assume that these can be obtained. We can then focus on the estimation of c_I.

The method used here is based on Greene (2003), §13.3 and §13.4.3. See also Wooldridge (2010), §10.5.3. In what follows, an illustrative distinction is made between an estimator (in the standard sense) and a predictor. The latter is obtained from data and estimators combined.

Consider the model in equation (3.3). For clarity, time-varying and time-invariant variables and coefficients are combined into X_{It} and γ respectively. The new model is:

$$\ln(w_{It}) = X_{It}\gamma + c_I + \epsilon_{It} \tag{3.4}$$

Now, assume $\hat{\gamma}$ is an unbiased and consistent (over N_I) **estimator** of γ. Define the estimation error $\hat{u}_{It}$ as:

$$\hat{u}_{It} \equiv \ln(w_{It}) - X_{It}\hat{\gamma}$$

Furthermore, define the linear **predictor** $\bar{u}_I$ as the average of the estimation error **over the sample period**:[9]

$$\bar{u}_I \equiv \frac{\sum_{t=1}^{T} \hat{u}_{It}}{T} = \overline{\ln(w_I)} - \bar{X}_I\hat{\gamma} \tag{3.5}$$

To evaluate the statistical properties of this predictor, replace the original model – equation (3.4) – into the above expression. After rearranging, the outcome is:

$$\bar{u}_I = \bar{X}_I\gamma - \bar{X}_I\hat{\gamma} + c_I + \frac{\sum_{t=1}^{T} \epsilon_{It}}{T}$$

Since the number of periods does not allow for large sample properties to hold, it is of interest to evaluate a small sample property of this predictor, namely its bias (or lack of). The focus is on the conditional expectation with respect to observed variables. This is:

$$E(\bar{u}_I|X_{It}) = \bar{X}_I\gamma - \bar{X}_I E(\hat{\gamma}|X_{It}) + E(c_I|X_{It}) + \frac{\sum_{t=1}^{T} E(\epsilon_{It}|X_{It})}{T}$$

[9] In the case of an unbalanced panel, number of periods is cluster-specific, i.e. T_I. Nevertheless, the result follows through.

By assumption, $\hat{\gamma}$ is an unbiased estimator of γ. Moreover, using the Law of Iterated Expectation, Appendix **??** shows that assumption A.1 implies $E(\epsilon_{It}|X_{It}) = 0$. In consequence:

$$E(\bar{u}_I|X_{It}) = E(c_I|X_{It})$$

This is, the average estimation error is an **unbiased** predictor of the conditional mean of c_I.

The conditional variance of this predictor is calculated from equation (3.5). This is:

$$Var(\hat{\nu}_I|X_{It}) = \frac{\Sigma_i}{T} + \bar{X}_I' V(\hat{\gamma}|X_{It})\bar{X}_I \tag{3.6}$$

where $V(\hat{\gamma}|X_{It})$ is the variance of the estimators. Recall that $\hat{\gamma}$ are consistent with respect to N_I. Although not shown here, they are also consistent with respect to T (e.g. Greene, 2003, p.288). In consequence, the predictor of (log) ability is consistent **with respect to** T. This is:

$$p\lim_{T\to\infty} \bar{u}_I = c_I$$

An alternative method to prove consistency of the predictor is to directly calculate its probability limit. This is:

$$p\lim_{T\to\infty} \bar{u}_I = p\lim_{T\to\infty}\left(\bar{X}_I\gamma\right) - p\lim_{T\to\infty}\left(\bar{X}_I\hat{\gamma}\right) + p\lim_{T\to\infty} c_I + p\lim_{T\to\infty}\left(\frac{\sum_{t=1}^{T}\epsilon_{It}}{T}\right)$$

As mentioned, it can be shown $\hat{\gamma}$ is a consistent estimator of γ over T. Therefore, $p\lim_{T\to\infty}\hat{\gamma} = \gamma$. Furthermore, from the Kolmogorov Law of Large Numbers, $p\lim_{T\to\infty}\left(\frac{\sum_{t=1}^{T}\epsilon_{It}}{T}\right) = E(\epsilon_{It})$, which is zero. The result is $\bar{u}_I \xrightarrow{P} c_I$, as before.

Finally, since $\bar{u}_I$ is an average of independent identically distributed variables, the standard ("Lindeberg–Lévy") version of the Central Limit Theorem ensures that:

$$\sqrt{T}\left(\bar{u}_I - c_I\right) \xrightarrow{d} N\left(0, \Sigma_I\right)$$

In consequence, the predictor, for each unit I, is $\sqrt{T}$-asymptotically normally distributed. Equivalently:

$$\bar{u}_{I|X_{It}} \sim N\left(c_I, \frac{\Sigma_I}{T} + \bar{X}_I' V(\hat{\gamma}|X_{It})\bar{X}_I\right)$$

Importantly, the normality of the predictor does not imply that c_I is normally distributed across I. It just means the predictor itself, being a random variable, and through the merits of the LLN, distributes normal in large samples. As such, this predictor does not impose any constraint on the distribution of c_I.

In summary, N_I is not enough for consistency; a long panel is. While this is not the case here, it is still true that the ability prediction is unbiased, as long as $\hat{\gamma}$ is also an unbiased estimator of γ. The latter is a positive result given the not so long panel available here. Naturally, the longer the panel, the more "accurate" this prediction is.

Finally, notice that c_I is predicted at the cluster level I. Thus, for switchers, there is a separate estimation for every occupation at which he or she worked during the sample.

Identification of "observed" ability

Having defined our estimator for c_I, we can go back to the theory. According to the model, this component represents the log of ability. This assumption can be thought from two angles, namely that ability is in fact a determinant of wages, and that no other unobservable factors beside ability are relevant for wages. The former should not be very controversial. It is rather ubiquitous in the Mincer literature that education is likely to be correlated with unobservables like ability. Hence the preference of Fixed Effects over Random Effects, or the use of instruments (as in the seminal work of Angrist and Keueger, 1991). Naturally, how ability impacts on wages is a separate issue, depending on a theoretical construct. The approach here is precisely that one, namely to start from a model, from where an econometric equation is derived. A more elaborated alternative would seek to build a fully structural model – popular since the seminal work by Keane and Wolpin (1997). This is a possibility to be explored in the future. For the interest of this chapter, the model used here is deemed to be enough.

The assumption that no other unobservable besides ability affects wages is surely a strong one. Two candidates that might be also omitted are firm heterogeneity and family background. Certainly, firms characteristics might directly affect pay. For instance, a manager could be paid more than another manager because the former belongs to a high-paying firm, and not because of higher ability. Similarly, a routine worker might be paid more in an assembly line at Airbus than at a beer factory, *ceteris paribus*. Unfortunately, the BHPS dataset does not allow us to explore these hypothesis, as the firm is not identified. The estimations do account for some of firm information like size, union status, etc, but cannot capture idiosyncratic firm effects.

But how high could this heterogeneity be? Starting perhaps with Abowd, Kramarz and Margolis (1999), a new literature has emerged using employee-firm datasets, which has proven to be an informative method to disentangle worker and firm effects on wages (for a recent survey, see Card et al., 2015). This methodology has been particular fruitful in evaluating recent changes in inequality, by decomposing between and within firm pay effects. But by doing so, they also evaluate the **level** of influence from both firm and worker's characteristics. For example, Card, Heining and Kline (2013) decompose the variance of West German workers' wage into several components, finding that the variance associated with unobservable person characteristics amounts to 0.127, whereas the variance associated to firm characteristics is 0.053, for the period 2002-2009 (Table IV, p.1000). For the case of the US, Barth et al. (2016) find the person and establishment effects being roughly similar (e.g. 0.189-196 in 2007; Table 3, p.S78) . In contrast, and for the UK, Burdett, Carrillo-Tudela and Coles (2016) find that firm heterogeneity explains between 29% and 41% of wage variance whereas worker unobserved effects account between 18% and 31% of it (Tables 3 and 4, p.37). Although the methodologies and population groups vary across these three studies, none of them consider the effect of wage rate effects. Since these depend on the occupation (partly related to firms) and not on worker's characteristics, the effect of heterogeneous firms net of occupation-specific wage effects is expected to be lower.

All in all, it is clear the effect of firms on worker's wages exists and it is not meaningless. The important question is whether this effect is captured by the **time-invariant, occupation-specific** component or not. As long as the mapping between firm and occupation is not 1-to-1, the contamination might not be too high. For example, if most of workers switch firms but not occupation, what c_I might capture is the average of the firms effects. However, if there is high degree of positive assortative matching (high ability workers going to high paying firms), the prediction of ability is partly confounding the effect of firm. Conversely, if workers who switch occupation do not switch firms, the effect of firm, constant across occupations, is taken away by the model's constant. Finally, as the literature just cited shows, there is an increasing importance of firm heterogeneity behind changes in variance of wages. It is possible that these variability in firm effects are partly captured by the wage rates and thus not passed through to c_I or ϵ_{It}.

In conclusion, the contamination from firm effects into ability effects depends on the degree of positive assortative matching and the pattern of firm-occupation linkages among switchers and non-switchers. Unfortunately, the dataset used here does not allow us to evaluate these factors.

Another factor that might play a role in c_I is family background. This is not a problem if and only if its effect on wages is only through workers' ability. Recall that the concept of ability here is relatively broad. It might include factors like confidence, OTHERS, that are known to affect productivity. Family background might be a problem if wages reflect access to networks, which give unequal access to high paying jobs, against the role of ability. Later a test is conducted where estimated c_I shows relatively little correlation with father's education. To the extent that the latter is a good predictor of the family environment, this suggests the omission of this factor might not be too worrisome.

Sorting model caveats

Several caveats can be deduced already from the above discussion, which hinder the capacity of the method to produce good estimates of ability. Central to all the above methodology is for ability in observed occupations to be well identified. This is essential not only for their own validity, but because the inference of counterfactual ability relies entirely on those estimations. Inconsistencies in the former invalidates the latter analysis. Naturally, achieving this consistency relies on (i) specifying the original model correctly, and (ii) obtaining a consistent estimation of it.

Nonetheless, even under consistent observed ability estimates, the full set of assumptions needed for predicting unobserved ability has to hold. There are many reasons why this might not be the case. First, the labour market is not strictly competitive. While most of workers cannot influence their pay (except through unions), very capable individuals – or "stars" – might be able to extract rents from firms (e.g. Choe, Tian and Yin, 2014). Competition is also hindered by reduced mobility across occupations. Switching is costly and might deter an otherwise profitable move (e.g. Cortes and Gallipoli, 2014). Job search has costs too and information is also not readily available to all. Moreover, many occupations require formal or informal qualifications which in effect create entry barriers, limiting the competition faced by already qualified workers. Additionally, experience might be occupation specific (e.g. Kambourov and Manovskii, 2009). As such, the prospect of occupational mobility might not be a realistic threat for many occupations.

Second, occupational decisions can be driven by both pecuniary and non-pecuniary motives (e.g. Farzin, 2009). For instance, some workers might have a stronger preference for certain areas or industries, or the quality of jobs might differ across firms. This is a common feature of many models in the literature. Rosen (1974, 1987) developed the so called Theory of Equalizing Differences. In his model, firms offer jobs that vary over certain attributes, and workers have preferences regarding wages and these attributes. Given standard competitive assumptions, the economy reaches an equilibrium where workers sort among jobs that better match their attributes. However, his model assume workers have the same ability level. Combining both unobserved heterogeneous preferences for attributes and unobserved heterogeneous ability renders identification of the latter very difficult.

Third, the labour market might not be in equilibrium. This is *per se* not a problem if shortages or excesses are reflected in wage rates. Yet, in markets where prices take time to adjust, these wages might not fully reflect market conditions, failing to reveal the incentives behind workers' occupational choices, and ultimately ability.

Last but not least, the derivation of unobserved ability is *ceteris paribus* in the strict sense. It assumes that workers evaluate occupations **holding other workers' decisions constant**. This might not be the case, for instance, when new opportunities arise for many workers. In this case, a precise evaluation should take into account the evaluation that other workers perform too, invalidating the *ceteris paribus* nature of the exercise. In other words, efficiency wages might not be truly exogenous to the evaluation exercise.

3.4 Estimation methods

So far nothing has been said about the statistical properties of the Data Generating Process underlying equation (3.3). To estimate this model, we need further assumptions. Since some of these assumptions might not hold, I try different estimation methods, each consistent under alternative assumptions. The first approach used is that of Random-Effects (RE). While this method assumes the existence of an unobserved effects, it imposes the strict restriction that it is uncorrelated with regressors. This is a rather unrealistic assumption, as ability is likely to be correlated with Education. In consequence, a Fixed-Effects (FE) method is also implemented, and standard tests are available to compare these two estimators.

An alternative to the FE estimator is the one proposed by Hausman and Taylor (1981) – HT henceforth. This method also allows for correlation between the unobserved component and ability, but instead of just estimating a demeaned model, in a second stage it uses the mean of exogenous regressors as "internal instruments" in a 2SLS procedure. The advantage of HT is that it provides consistent estimates for the time-invariant coefficients without requiring long panels (as FE does). This comes at the cost of extra orthogonality assumptions, not needed in FE.

Moreover, education can also be endogenous with respect to the idiosyncratic error, possibly due to simultaneity bias. This is a natural consequence of education being itself a choice that depends on expected level of wages. Under this scenario, RE, FE and HT are inconsistent. To account for this possibility, an extra method – denoted FE-2SLS – is applied. This is also an instrumental variables approach, but within the FE framework. Unlike the HT estimator, FE-2SLS uses external instruments (i.e. variables that are not included in the model). Several tests are available to evaluate the relevance and validity of these instruments, and to test whether education is indeed endogenous.

These four estimation methods are described in more detail below.

Random Effects

The random effects estimator is a Feasible Generalised Least Squares (FGLS), based on an estimation of a weighting matrix which adjust for serial correlation and heteroskedasticity of the composite error $c_I + \epsilon_{It}$.

Consider the following three assumptions:

A.1: Strict exogeneity

$$E(\epsilon_{Is}|\boldsymbol{x}_{It}, \boldsymbol{z}_I, \omega_{jt}, c_I) = 0, \quad \forall s, t = 1, \dots, T$$

A.2: Ability uncorrelated with regressors and wage rates

$$E(c_I|\boldsymbol{x}_{It}, \boldsymbol{z}_I, \omega_{jt}) = E(c_I)$$

Assumption A.1 implies that, after controlling for a worker's characteristics, wage rates, and ability, all movements in wages are purely random. This rules out lagged variables in the equation, and phenomena like workers learning their own ability – revealed through the error term. It also assumes no endogeneity due to simultaneity bias, or omitted variable bias. Assumption A.2 means that ability cannot be correlated to any regressor (e.g. education). It also rules out non-competitive effects, like a worker affecting the market wage rate. Given these two assumptions, RE provides unbiased and consistent estimates of α, β, and the occupation-year effects $\ln(\omega_{jt})$ (Greene, 2003, §13.4). It can be shown that the variance of the RE estimator degenerates to zero as either N_I or T go to infinity.

Fixed Effects

Fixed effect methods are needed when assumption A.2 fails. This is the case if ability is correlated with, for example, education. This correlation is most likely to be positive, as ability and education are complements in the generation of human capital, which in the end is what might define a worker's productivity at work. In consequence, if this correlation is not accounted for, estimated returns to education can be biased upwards. As shown later, this is actually the case for the sample under study.

The FE estimator is based on a demeaned version of equation (3.3). This is, for each variable, including wage, its cluster mean is subtracted. Since the unobserved component is constant within clusters, it is eliminated. In the end the FE estimator is just an OLS estimator over the demeaned equation. Since there is no unobserved component in the model any more, for unbiased and consistent estimation under FE, only assumption A.1 is required. Again, it can be shown that the variance of the FE estimator is degenerate for either N_I or T going to infinity (e.g. Greene, 2003, §13.3, and equation (13-8) in particular).

Predicting time-invariant variables under FE

Recall that to obtain an unbiased and consistent predictor of ability, we need an unbiased and consistent predictor of all model parameters, including time-invariant ones. Yet, FE does not provide an estimation of the latter. As suggested by Kripfganz and Schwarz (2015), a solution is to regress the estimated errors from the FE equation – which contains all time-invariant observables and unobservables – against the time-invariant variables not estimated by FE, a method based on Hausman and Taylor (1981). As it turns out, this solution allow us to recover these parameter without strong extra assumptions. Here I adapt Kripfganz and Schwarz (2015) method to our particular setting.[10]

Without loss of generality, consider equation (3.3) without the wage rate variables:

$$\ln(w_{It}) = x_{It}\alpha + z_I\beta + c_I + \epsilon_{It} \tag{3.7}$$

FE estimates α, but it **does not** estimate β. Define the estimation error in the FE equation, $\hat{v}_{It}$ as:

$$\hat{v}_{It} \equiv \ln(w_{It}) - x_{It}\hat{\alpha}_{FE}$$

[10]Here the time-invariant regressors are assumed to be exogenous, and there is no lagged dependent variable. Additional, Kripfganz and Schwarz (2015) use a standard clustering, whereas here I use occupation-spell clustering.

The linear predictor $\bar{v}_I$ is:

$$\bar{v}_I \equiv \frac{\sum_{t=1}^{T} \hat{u}_{It}}{T} = \overline{\ln(w_I)} - \bar{x}_I \hat{\alpha}_{FE}$$

In order to study the properties of this predictor, replace equation (3.7) into the above expression. This yields:

$$\bar{v}_I = \bar{x}_I \alpha - \bar{x}_I \hat{\alpha}_{FE} + z_I \beta + c_I + \frac{\sum_{t=1}^{T} \epsilon_{It}}{T} \tag{3.8}$$

The conditional expectation of this estimator with respect to observed variables is:

$$E(\bar{v}_I | x_I, z_I) = \bar{x}_I \alpha - \bar{x}_I E(\hat{\alpha}_{FE} | x_I, z_I) + z_I \beta + E(c_I) + \frac{\sum_{t=1}^{T} E(\epsilon_{It} | x_I, z_I)}{T}$$

Appendix **??** uses the Law of Iterated Expectations to show that assumption A.1 implies $E(\epsilon_{It} | x_I, z_I) = 0$. The latter is also sufficient for $E(\hat{\alpha}_{FE} | x_I, z_I) = \alpha$. In consequence:

$$E(\bar{v}_I | x_I, z_I) = z_I \beta + E(c_I)$$

This is, our predictor is an **unbiased** estimator of the time-invariant components of the model, a desirable property in small samples.

Regarding consistency, the probability limit of this predictor is:

$$p \lim_{T \to \infty} \bar{v}_I = p \lim_{T \to \infty} (\bar{x}_I \alpha) - p \lim_{T \to \infty} (\bar{x}_I \hat{\alpha}_{FE}) + p \lim_{T \to \infty} z_I \beta + p \lim_{T \to \infty} c_I + p \lim_{T \to \infty} \left(\frac{\sum_{t=1}^{T} \epsilon_{It}}{T} \right)$$

As stated before, $\hat{\alpha}_{FE}$ is a consistent estimator of α, whereas the Kolmogorov Law of Large Numbers ensures the error term above converges to its mean, which is zero. Therefore:

$$p \lim_{T \to \infty} \bar{v}_I = z_I \beta + c_I$$

This is, our predictor is a **consistent** estimator of the time-invariant components of the model, and unbiased in small samples. Furthermore, just as before, the "Lindeberg–Lévy" version of the Central Limit Theorem implies:

$$\sqrt{T}(\bar{v}_I - z_I \beta - c_I) \xrightarrow{d} N(0, \Sigma_I)$$

which in turns leads to:

$$\bar{v}_I | x_I, z_I \sim N\left(z_I \beta + c_I, \frac{\Sigma_i}{T} \right) \tag{3.9}$$

In words, $\bar{v}_I$ is *as if* normally distributed, with conditional mean $z_I\beta + c_I$.

Now, in practice we observe z_I but do not observe c_I. Yet, using the conditional mean just derived, we can proceed to estimate β and c_I. More precisely, consider the following model:

$$\bar{v}_I = z_I\beta + \xi_I \ , \qquad I = 1, \ldots, N_I \tag{3.10}$$

where $\xi_I \equiv c_I + \bar{x}_I(\alpha - \hat{\alpha}_{FE}) + \frac{\sum_{t=1}^{T}\epsilon_{It}}{T}$, according to equation (3.8). Here again we can see that $\bar{v}_I$ distributes as in equation (3.9). In order to obtain a consistent estimation of β (conditional on $\{x_I, z_I\}$), we need to assume that z_I is exogenous in the above equation, i.e. $E(C_I|z_I) = 0$. Since our main concern is that ability is correlated with education (part of x_{It}), this assumption might not be very strong. Then, it follows that β can be consistently estimated with OLS.[11] This is:[12]

$$p \lim_{N_I \to \infty} \hat{\beta}_{OLS} = \beta$$

where $\hat{\beta}_{OLS} = (Z'Z)^{-1}\left(Z'\bar{V}\right)$, and Z, $\bar{V}$ represent data as matrices. It can also be shown that $\hat{\beta}_{OLS}$ is an unbiased estimator of β.

Although this is enough to estimate ability using the method presented in Section 3.3, it turns out the **mean** estimated error of equation (3.10) is also an unbiased and consistent predictor of c_I. More precisely, define this predictor as:

$$\hat{\xi}_I = \bar{v}_I - z_I\hat{\beta}_{OLS}$$

It is clear then that $E(\hat{\xi}_I|x_{It}, z_I) = c_I$.

Moreover, since the asymptotic variance ("Avar") of $\bar{v}_I$ is $\frac{\Sigma_I}{T}$, it also follows that:

$$Avar(\hat{\xi}_I|x_{It}, z_I) = \frac{\Sigma_I}{T} + z_I' Avar(\hat{\beta}_{OLS}|x_{It}, z_I)z_I$$

where $Avar(\hat{\beta}_{OLS})$ is calculated as usual for OLS. Since $\hat{\beta}_{OLS}$ is $\sqrt{N_I}$-asymptotically normally distributed, and $p \lim_{N_I \to \infty} \frac{\Sigma_I}{T} = \frac{\sigma_i}{T}$ (again, see Greene, 2003, p.288), the

[11] Recall that the clustering at I induces cross-sectional dependence, which violate Gauss-Markov assumptions. Still, it was argued before that the extent of CSD is bound to be very low.

[12] The proof of this is straightforward, and not shown here. The result derives from the orthogonality assumption, which implies $p \lim N_I \to \infty \left(\frac{Z'c}{N_I}\right) = 0$

asymptotic variance of $\bar{v}_I$ converges to $\frac{\sigma_i}{T}$ when $N_I \to \infty$. This result is in line with the general consistency condition highlighted in the identification method. Namely, that ability can be consistently estimated only when $T \to \infty$.

Before moving forward, some clarifications concerning the estimation of this second stage regression are necessary. First, since the dependent variable is a generated variable, $\hat{\beta}_{OLS}$ standard errors are wrong and subsequent test statistics are invalid (Wooldridge, 2010, §6.1.1). Robust standard errors can be obtained using bootstrapping methods. Second, switchers have more than one observation I in the above sample. If their ability levels are correlated among tasks (for example, if there is high degree of absolute advantage), the error terms in equation (3.10) are correlated among observations I corresponding to the same individual i. This can be corrected by performing the bootstrap re-sampling above at the cluster level i.[13] Third, as already mentioned, the quality of this second stage estimation hinges crucially on the quality of the predictor, which in turn depends on the **length** of the panel. In short panels this is likely to be a noisy estimate, yet unbiased. While nothing can be done to improve this here, bootstrapping methods can help to quantify the extent of this noise.

Hausman-Taylor

The third estimator used was proposed by Hausman and Taylor (1981) – HT henceforth. Just like the FE estimator, HT also allows for correlation between ability and explanatory variables. In particular, HT divides regressors into four categories, depending on whether variables are time-invariant or not, and whether they are correlated with the unobserved component or not. In our case, education is time-varying, which means there is no fourth category of time-invariant endogenous regressors.

To see how this method works, consider equation (3.3), where endogenous regressors are separated from exogenous ones. The model is:

$$\ln(w_{It}) = x^*_{It}\alpha^* + educ_{It}\alpha_{educ} + z_I\beta + \sum_{j\times T} D_{It}\ln(\omega_{jt}) + c_I + \epsilon_{It} \tag{3.11}$$

[13] In effect, the model in equation (3.10) can also be seen as a pooled cross-section model, by replacing I with i, j. The dependent variable and the error term vary over i and j, whereas the time-invariant variables only vary across i. Clustering over i is then natural in order to obtain robust standard errors.

where x^*_{It} represents all (exogenous) regressors but education. HT method is as follows. First, a FE estimation of equation (3.11) is carried out, from where $\hat{\alpha}$ is obtained. Here, the estimated errors are computed as usual. Second, these errors are regressed against the time-invariant variables z_I, using themselves and demeaned x^*_{It} as instruments, in a 2SLS procedure. Hausman and Taylor, 1981 shows these instruments provide consistent estimated of β. However, these estimates are inefficient (e.g. Greene, 2003, §13.5). Third, from these estimates, an estimation of the variance of both c_I and ϵ_{It} is computed. These are then used to calculate a weighting factor $\hat{\theta}_I$, for each cluster I, which are then used to transform the original model in equation (3.11), as is also done under RE. Fourth, this transformed model is estimated with 2SLS, using z_I and demeaned x^*_{It} as instruments. Since the data has been transformed using a consistent estimation of the model's error structure, the latter estimation is not only consistent but also efficient. Because of this final step, the HT estimator is in essence a RE-2SLS estimator.[14]

Since our concern is only with education – a time-varying variable – using HT does not capture it's full benefits. In consequence, there is little difference between FE and HT estimations, as shown later. The crucial difference however is that HT estimates the time-invariant variables jointly with the time-varying ones – as RE does, a property which is missing in FE. Thus, it is still informative to carry on this method. Additionally, the wage rate can be estimated for all task-years (except one if constant is added), in contrast with FE.

FE-2SLS

The three estimation methods proposed so far assume education is uncorrelated with the idiosyncratic error. FE and HT estimations allow for correlation of education and ability, but this might not capture all the potential endogeneity. For example, in many ways education is a choice, based on expectation about future wages. In consequence, there could be reverse causality, even after accounting for ability. If this is the case, assumption A.1 is violated, and RE, FE and HT estimators are biased and inconsistent. The solution is either to build a structural model that takes into account educational

[14] Yet, unlike the FE-2SLS estimator, HT **does not** allow endogeneity with respect to the idiosyncratic error.

choices (for example, Heckman, Humphries and Veramendi, 2016), or use external instruments to isolate the effect of education on wages from other confounding factors. Here, I use the second approach. This fourth estimator – called FE-2SLS – runs a standard 2SLS on top of a FE. The instruments used are the education level of the individual's partner (if any) and that of the father.[15] One of the assumptions needed for FE-2SLS to provide consistent estimates is strict exogeneity (except for the endogenous variable, of course). The modified version of A.1 is as follows:

A.1b: Strict exogeneity

$$E(\epsilon_{Is}|\boldsymbol{x}^*_{It}, \boldsymbol{z}_I, \boldsymbol{q}_I, c_I) = 0, \quad \forall s, t = 1, \ldots, T$$

where q_I are the external instruments, and x^*_{It} does not include education. Importantly, this assumption allows for arbitrary correlation between the instruments and the unobserved component. This is likely to be the case, as own ability and partner or parents' education could be positively related. The other assumptions needed are standard in 2SLS. Of special relevance is the validity and relevance of the instruments used. Since there are two instruments, these conditions can be tested.

Occupation specific ability test

A central tenet in this whole analysis is the assumption that workers' ability is occupation specific. This is, that η_{ij} varies across j. If this is not the case, estimation methods presented earlier could be quite inefficient. This is because, instead of using the correct cluster (i), they are performed under a reduced cluster ($I(i, j)$). This does not make a difference for never switchers but it matters for switchers. This is specially acute for methods that rely **only** in within cluster variability (FE and FE-2SLS). Of course, an even more dramatic consequence of rejecting this approach is that the theory under study has to be rejected.

On the contrary, if there is compelling evidence that ability does vary across occupations, the implications for the literature are quite fundamental. In effect, using the wrong

[15] There is very high correlation between both parents' education level, so only one parent can be used at the time. Where father's education is missing, the mother's is used instead.

clustering (i) would mean that FE methods applied to Mincer equations (in my view, the standard) do not eliminate unobserved heterogeneity. This is because, for switchers, $\frac{\sum_t \eta_{ij}(t)}{T} \neq \eta_{ij}$. Therefore the demeaned effect, equivalent to $\eta_{ij} - \frac{\sum_t \eta_{ij}(t)}{T}$, and always zero with correct clustering, is not zero. The extent of the problem depends on how many in the sample under study are switchers. Notice that wrong clustering also invalidates the Hausman test (which assumes both models are consistent under the null), widely used to compare RE and FE models. Interestingly, there is already evidence in the literature that ability varies across dimensions. The work of Lemieux (1998) and Imbert (2013) mentioned in Section 3.2 point to the same conclusion. Yet, they do not highlight the relevance of this result for common Mincer equations, as I do here.

Given the relevance of this issue, this assumption is put to the test. For this, I design a test based on Hausman test's underlying methodology. Consider the following two models:[16]

$$\text{Model A:} \quad y_{it} = X_{it}\beta + \chi_i + error \tag{3.12}$$

$$\text{Model B:} \quad y_{it} = X_{it}\beta + \zeta_{ij} + error \tag{3.13}$$

where χ_i and ζ_{ij} are an unobserved ability component, correlated with some elements in X_{it}. In the first model ability is general whereas in the second model it is occupation specific. Now, consider two estimation methodologies:

- Standard FE: a fixed effects estimation for a panel defined at the i cluster. This is, the fixed effect estimator differentiates the data using means **across individuals**, as in a standard FE estimation. This method eliminates the unobserved effect in Model A, but in Model B it only eliminates the unobserved effect for individuals who **do not switch occupations** in the sample. In consequence, in a sample with switchers, a standard FE cannot fully eliminate the unobserved component. This leads to inconsistent estimations because the unobserved effect is still correlated with observables. The estimation of Model A however is consistent and efficient.

- Occupation-spell FE: a fixed effects estimation for a panel defined at the $I(i,j)$ level. This is, an individual at different occupations represent different units over

[16] The omission of the time-varying wage rate factors is inconsequential to the results shown here.

which the mean-differentiation is carried out. While this method still eliminates the unobserved component in Model A, it also does so for both switchers and non-switchers in Model B. In consequence, this method provides consistent estimation for both models.

Now, we can see how a Hausman test fits under this scheme. Consider the null hypothesis that Model A is correct against the alternative that Model B is correct. As argued before, the occupation-spell FE estimator is consistent under **both** the null and the alternative, whereas the standard FE is consistent **and** efficient under the null, but it is inconsistent otherwise. This is the exact setup of the Hausman test. As such, this test can reveal information about whether Model A or Model B are accurate descriptions of unobserved effects.[17]

3.5 Data

Sample

The data comes from the British Household Panel Survey (BHPS), a longitudinal dataset covering Great Britain from 1991 to 2008. The sample starts from a random draw of households in the year 1991, which are followed thereafter. The sample is extended a few times. For example, in 1999, an extra sample from Scotland and Wales was obtained, to increase these countries' representativeness. Similarly, in 2001 Northern Ireland was added, hence extending the survey to the whole of the United Kingdom. Besides these extensions, the sample fluctuates for natural reasons. Some households and individuals exit the sample through death or attrition, whereas other enter through marriage and birth. From this dataset I select all those individuals for which wage and control variables (described in the next section) are available in a given year. Those with missing data in a particular year are excluded from the sample in that year, although some imputation was implemented when reasonably unambiguous (see Section G.2). As such, the panel is unbalanced. I dropped observations in the top 0.5% of the real

[17] To the best of my knowledge, I am not aware of any paper that adapts the Hausman test to evaluate the underlying dimensions of unobserved heterogeneity, as in here. If this is the case, the test presented here is a novel contribution to the literature.

income scale (corresponding to those with a real monthly wage above $5,600). This is because their pay might be influenced by factors unlikely to be captured by worker or job characteristics alone (including ability), as highlighted by the recently booming CEO pay-performance literature.[18]. As it turns out, 87% of these top income earners are classified in the abstract task (corresponding mainly to managerial and professional positions), in contrast with the 34% across the whole sample. Similarly, I dropped those at the bottom 0.5% (corresponding to those with a real monthly wage below $70), as these could be highly influenced by measurement errors or non-market phenomenon.[19] Additionally, those working less than 10 hours a week and more than 60 hours a week are dropped under the suspicion that these abnormal working hours might be less likely to reflect market conditions. These amount to around 2% of total observations. The final adjustment to the sample was the exclusion of members of armed forces, and those self reported as farm managers (see section 3.6), which combined make up less than 0.1% of the sample. Overall, observations dropped that could have been otherwise included in the model add up to around 3% of the total valid sample.

Some characteristics of the final sample are shown in Table 3.1. Male and females are more or less evenly represented in the sample. Of particular interest is the race composition. Only 2.7% of observations correspond to non-white individuals. For reference, the proportion of non-white among the UK population was 7% and 9% according to the 1991 and 2002 censuses, respectively.[20] Moreover, Panel (a) in Figure 3.5 shows the distribution of the number of years individuals are present in the sample. Around 4% of individuals are present during all years. The peak at 8 years and later drop reflects the late inclusion of Northern Ireland into the sample. This figure shows that the panel is very unbalanced. There is considerable entry to and exit from the sample, either temporary (e.g. changing employment status) or permanent (e.g. attrition). This suggest selection bias might be a serious problem, and has to be properly accounted for. Finally, Panel (b) in Figure 3.5 shows the evolution of sample size over the years, confirming the prevalence of entry and exit.

[18] For example (Gabaix and Landier, 2008) and (Gabaix, Landier and Sauvagnat, 2014) explain most of the increase in CEO pay in the US in the last thirty years based on the increase in the market capitalisation of the top firms. In our case, the variable reflecting the size of the firm is measured in terms of the number of employees, a poor indicator of market capitalisation (Bryan, 2007).

[19] Many of these individuals reported a monthly pay of $0.1.

[20] Yet, censuses include non-working adults, children and elderly, whereas this dataset does not.

Table 3.1: Description of sample, 1991-2008

	N	**Percentage**
Total observations	90,780	100%
Male	44,076	48.6%
Female	46,704	51.4%
White	88,339	97.3%
Non-white	2,441	2.7%
Full-time	73,746	81.2%
Part-time	17,034	18.8%
Public sector	28,870	31.8%
Private sector	61,910	68.2%
Union member	46,643	51.4%
Not member	44,137	48.6%
Total individuals	15,843	100%
Male	7,670	48.4%
Female	8,173	51.6%

Variables

The dependent variable used in this study is the logarithm of the real monthly wage. The original variable taken from the BHPS – paygu – is defined as the latest gross wage the individual had received at the moment of the interview. This is a nominal variable, and includes only labour earnings. It was deflated using the Retail Price Index, compiled by the Office of National Statistics. The base period is the third quarter of 2000. Notice that the use of real wage instead of nominal wage means the estimated wage rates components for each task-year interaction are also in real terms.

The regressors included in the model are standard in the Mincer literature. The first key variable included is years of education. This is calculated by subtracting 16 (age at which secondary school finishes) to the self-reported age at which further education finished. As this is a longitudinal survey, the latter can change over time, which means years of education can also change. However, in the wage regression this is only relevant if students are also working, or workers rejoin further education later in their lives. For the majority of the population, education is time-invariant. An alternative method to

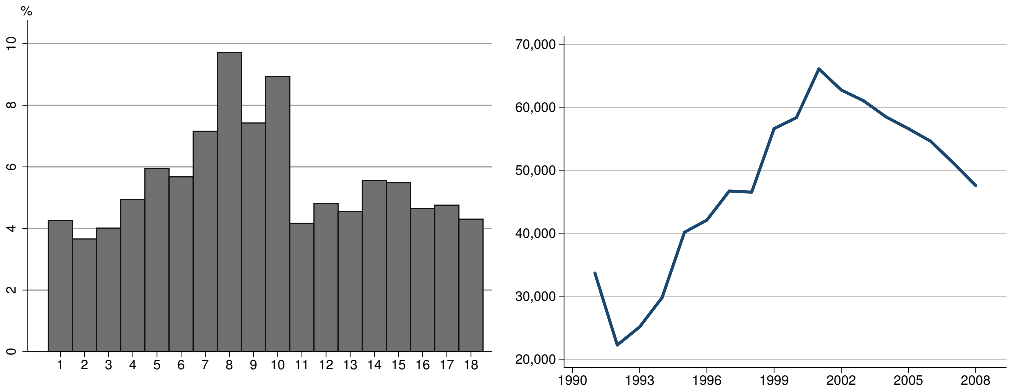

(a) Average number of years in sample among individuals

(b) Sample size over the period

Figure 3.5: Some sample characteristics

account for education is to introduce dummies for each educational category. Although this is not feasible for HT and FE-2SLS estimators, it is also implemented in RE and FE. The benefit is that it allows for arbitrary differences in the effect of each degree types. Additionally, in the later two models, education is allowed to vary across gender.

The second essential regressor is experience. This is normally calculated as current age minus graduation age. Here I pursue a more sophisticated method. I dynamically calculate experience using the in-sample job market status of the individual, in order to adjust experience gained from full-time and part-time jobs. In my scale, experience gained in full-time job for a year adds 1 "full-year equivalent" year of experience; experience gained in a part-time job alone add up 0.5 "full-year equivalent" years of experience; and experience gained while studying and working part-time is 0.25.[21] Those outside employment are assumed to gain no experience. Table 3.2 presents an example of this method in detail. The individual moves from full-time education into full-time work, spending some years as a part-time worker and student, and a year unemployed. The third column calculates this worker's "stock" of experience at the **beginning** of the period, based on previous years' accumulation. The highlighted rows indicate the individual is part of the estimation sample.

For those who join the sample having already finished their education, the standard method is used. It assumes they have been in full-time employment since graduation,

[21] Notice that the effect of more education is already accounted for in the education variable.

Table 3.2: Example showing how experience is calculated

Age	Status	Experience
⋮	⋮	⋮
17	Full-time study	0
18	Full-time study	0
19	Full-time work	0
20	Full-time work	1
21	Part-time study and work	2
22	Part-time study and work	2.25
23	Part-time study and work	2.50
24	Full-time work	2.75
25	Full-time work	3.75
26	Full-time work	4.75
27	Unemployed	5.75
28	Part-time work	5.75
29	Part-time work	6.25
⋮	⋮	⋮

Note: highlighted rows indicate the individual is part of the estimation sample

due to the lack of a better alternative. As usual, experience squared is also introduced in the model. Moreover, Lemieux (2006) suggest fitting a quartic polynomial for experience to better account for those inexperienced just entering the labour market. This is also implemented. There is also a distinction by gender.

Notice that the definition of experience used here is compatible with the concept of generic skills, opposed to skills specific to the firm or to the current employer. To account for specific skills, some studies also introduce tenure. While BHPS does contain a variable reporting consecutive days of work with current employer, this has huge measurement errors. Results are quite often inconsistent (e.g. contradictory between successive interviews), or implausible (e,g, some report around 100 years of tenure). For this reason, this variable is not used.

The model includes several individual characteristics not related to the job like gender, race and marital status. Race is defined as white and non-white. Marital status, based on the “actual” marital status as opposed to the “legal” marital status, is categorised as never married, couples/married, and divorced/widowed. Of these, marital status is the only time-varying.

The rest of the regressors are job attributes related to the firm. There is a public sector dummy, an indicator for full-time or part-time contract, the number of hours usually worked, a dummy for self-employment, and whether there is a trade union at workplace. Also, I include the firm's size in terms of number of employees (coded in five categories, ranging from less than 10 employees to more than 500 employees), and the industry where the firm operates. The latter is defined in 12 categories, taken from the SIC92 standard. Because some of them had very few observations (inducing collinearity problems), I had to merge some of them.[22] Finally, the model includes the region where the individual lives, coded in 12 groups across the UK. Again, some had to be merged to avoid collinearity issues.[23] Since there is no variable defining the location of the firm, this is used as a proxy. Some of these regressors are also differentiated by gender, based on pre-tests.

3.6 Tasks and Occupations

Classification of tasks

The theoretical model upon which the econometric model is based assumes there are three tasks in this economy: manual, routine, and abstract. This classification has become relatively common within the *task* literature, and it is implicit in the seminal contribution by Autor, Levy and Murnane (2003), from where this literature emerged. In effect, the first distinction they make between tasks is that of routine and non-routine tasks. As they say, "a task is 'routine' if it can be accomplished by machines following explicit programmed rules" (p.1283). Thus, these known rules can be implemented

[22] The final list of industries included is Agriculture, hunting, forestry, fishing, mining and quarrying (as a single industry); Manufacturing; Construction, Electricity, Gas and Water Supply; Wholesale and Retail Trade; Hotels and Restaurants; Transport, Storage and Communication; Financial Intermediation; Real Estate, Renting and Business Activities; Education; Health and Social Work; Other Community, Social and Personal Service Activities; and Public administration.

[23] Regions include London, South East, South West, East Anglia, East Midlands, West Midlands, Manchester and Liverpool, Rest of North West, Yorkshire and the Humber, North East, Wales, and Scotland and Northern Ireland.

into codes and computers. On the contrary, non-routine tasks cannot because "the procedures for accomplishing them are not well understood" (p.1283).

Secondly, they classify tasks as manual and cognitive (they also denominate the latter as analytical and interactive). As such, they produce a 2x2 categorisation of tasks: manual routine and non-routine, and cognitive routine and non-routine (see Table I in original paper). Then, they argue that the substitution effect of computerisation is very strong in routine tasks, regardless of they are manual or cognitive. Additionally, they claim that the effect on computerisation if also differentiated between manual and cognitive tasks, being relatively indifferent to the former but complementary to the latter. It is this important distinction on the role of technology on tasks which has motivated the literature to bundle together "manual routine" and "cognitive routine" tasks – simply calling them routine, and to differentiate between "manual non-routine" and "cognitive non-routine" (which some authors also call abstract). This final categorisation has more or less prevailed in the literature ever since.[24]

Naturally, this simple, three tasks approach to occupations and technology in general is not without critiques. One of the main problems refer to the lack of a precise classification of tasks.[25] As said, Autor, Levy and Murnane (2003) classified tasks based on two dimensions – "routineness" and manual/cognitive. Blinder (2009) adds offshorability as another dimension. This is clearly interesting in itself, since offshoring of tasks has been recognised as an important driver of recent employment trends (e.g. Fortin and Lemieux (2016); Oldenski, 2014). While cognitive tasks are less prone to offshoring, both routine and non-routine tasks are. Then, the correct understanding of technology (which enhances off-shoring through, e.g. lower communication costs) on employment reallocation might not be solely captured by the standard three tasks.

Another critique comes from Susskind (2016). He notice that the standard classification of tasks as routine and non-routine is based on a rather old-fashion conception of

[24]After this classification has been established, the introduction of task-biased technical change is straightforward. As technology is more prone to substitute workers in routine-intensive occupations, these workers are replaced by computers, machines, or robots, fostering employment away from these occupations. The link with job polarisation is also straightforward. For example, see Autor and Dorn (2013).

[25]See Autor and Handel (2013) for (in my opinion) one of the most precise classification of tasks available in the whole literature. It is based on O*NET.

technology. He argues that machines are increasingly performing tasks that are classified as non-routine, effectively destroying the original Autor, Levy and Murnane (2003) boundary. He proposes a more graded approach, based on the degree of "routinisability" of tasks. There is clearly some truth in Susskind's approach, in particular with respect to the automation of an increasing number of tasks normally considered non-routine. Yet, he is much more focused on the future of jobs and recent breakthrough in automation, whereas the perspective here (and in great part of the literature) is mainly on understanding the past. The strong evidence of past polarisation using standard routine versus non-routine categories indicates that it still remains informative when studying existing data.[26] As time goes by, it is perhaps pertinent to consider Susskind (2016) argument more seriously. Finally, to give the last word to one of the leading authors of this theory, see Autor (2013) for a more recent apology of this approach to labour markets.

Task-occupation mapping

Ideally, we would like our longitudinal and employment datasets to include a variable called "task" – with categories manual, routine, and abstract, over which workers are classified. This is however not the case. As said, tasks are a (recent) theoretical construct, for which no official standards exists (opposite to occupations or industries, where many standards exist). In fact, there is no agreement in the literature about how would such standard be defined. All this means we need to construct that "task" variable ourselves.

The model in this chapter assumes workers have an specific ability for each task (manual, routine, and abstract). Moreover, workers allocate into jobs where they perform a single task. This means the model assumes a 1-to-1 mapping between tasks and occupations. In consequence, one way to produce such "task" variable is to map occupations (which are normally included in surveys) into tasks. This is, for every occupational code available in a given survey, we associate one of the three tasks. This is the approach taken here

[26] Autor, Levy and Murnane (2003) make a similar argument. They do recognise the development of technologies like neural networks, which are allowing computers to e.g. recognise patterns (a remarkable advancement in visual technology, and probably an essential element in the incoming driveless car revolution.). Still, they claim that these should have not had a significant impact on the "computer revolution" of the last decades, which they set to estimate. See footnote 6 in their paper.

to produce our task variable, which will then link datasets of workers characteristics with our theoretical task construct.

Now, this 1-to-1 mapping is not straightforward. Usually, jobs involve multiple tasks. For example, Bisello, 2013 evaluates the task content of different occupation groups using the UK Skills and Employment Survey. This is a survey which includes very detailed information about the type of activities workers perform in their jobs. By linking these activities – which she classifies as analytical, interpersonal, routine, and manual – with workers reported occupation codes, she is able to calculate the extent each occupation uses each of these four tasks. The results are quite revealing.[27] For most of the occupations there is a strong correlation between analytical and interpersonal tasks (providing further support to the unified "cognitive" or abstract construct used in the literature). However, there is also relatively high correlation between manual and routine (see Table 5 in original paper). This means classifying occupations as abstract is less problematic than classifying them as manual or routine.[28] For instance, Bisello finds that the content of Agricultural and fishery labourers to be relatively equal in both manual and routine. The only way to move from these results to our 1-to-1 mapping is by introducing some theoretical intuition or knowledge. For example, if we were to classify Agricultural and fishery labourers, they would be manual workers because of the lower likelihood of these tasks to be prone to automation (opposed to, e.g. Customer services clerks, who also present relatively high levels of manual and routine, and which would be classified as the latter).

An alternative approach to the aforementioned 1-to-1 mapping is for each occupations to be a composition of tasks. This composition might be unique to each occupation. This is of course an interesting approach, which can lead to a precise analysis of tasks prices and the changes in tasks composition within occupations (e.g. Autor and Handel, 2013; Firpo, Fortin and Lemieux, 2011; Fortin and Lemieux, 2016). This however imposes considerable challenges for the estimation of ability from unobserved heterogeneity, which

[27] Sadly, she only report results for occupation codes aggregated at the two-digit level. The original data includes codes at the four-digit level. This is partly due to sample sizes problems (i.e. that at the most disaggregate level, many occupations only have a few workers).

[28] Following Autor, Levy and Murnane (2003) 2x2 categorisation of tasks, this might imply that the number of occupations that involve manual **and** routine tasks is much greater than the number of occupations that involve manual and non-routine tasks. If the opposite where true, the correlation between routine and manual would be lower.

is the central goal here. In effect, the method used here assumes that the unobserved heterogeneity component stands for unobserved ability, which enables its identification. Using this same unobserved heterogeneity to identify three ability components (one for each task) seems impossible.

The mapping used in this chapter

As discussed above, in order to put the model to the test, it is necessary to map each occupation code available into one of the three task used here. The BHPS provides occupational information of workers based on three "standards", namely ISCO88, SOC90, and SOC00. The ISCO (International Standard Classification of Occupations) is an international standard defined by the International Labor Organisation, whereas the SOC (Standard of Classification of Occupations) is particular of the UK, administered by the Office of National Statistics. Each classification is composed of major, sub-major, minor and unit groups, with 1, 2, 3, and 4 digits categories respectively.[29] Yet, each standard uses their own classification and numbering, and they cannot be directly compared. Still, worker's occupation are usually classified with two standards, meaning some robustness check can be conducted between mappings.

The mapping is based on a careful study of the related literature. Some of these use different standards, so the information they provide is not equivalent. For example, Bisello (2013) evaluates tasks content using ISCO88. This is particularly interesting because of the specific task content analysis of each occupational groups, albeit at the two-digit level (here the focus is on four-digit codes). Holmes and Mayhew (2015) uses SOC00 to map occupations into tasks, up to the four-digit level. Cortes (2016) also uses the highest disaggregation possible, using the COC standard for the US. The approach taken here is to adopt those classification that are highly consistent between them. This is the case of several categories like major groups 1 and 2 of the SOC and ISCO codes, which compromise managers and professional occupations, and are regarded by cited studies as *abstract* tasks. There is also an unambiguous mapping between clerical occupations, skilled trades, sales and customer services occupations, process, plant and

[29] For example, in SOC2000 an economist is classified as a unit group of 2322, minor of group of 232, sub-major group of 23, and major group of 2.

machine operators and *routine* tasks. For some categories, there is relative consensus at the aggregate but some disagreement at the finer level – like "Associate professional and technical occupations" (major group 3 in each standard). This include, for example, Engineering technicians (classified as *abstract*), and Train drivers (*routine*). The criteria to disentangle contradictions was two-fold. On the one hand, information about the precise activities or tasks performed by a given occupation was studied, based on official documentation.[30] On the other hand, the categorisation that is of interest in our case relates to the role of technology leading to automation. Some tasks with (at least until recently) low risk of automation – like low skill services (e.g. hairdressing) and low skill production occupations (e.g. farming), are regarded as manual, even though they might have a high routine content too. All in all, the degree of discrepancies between standards was never major. Among the highest disagreement of all were nurses and farm managers. The former were regarded as abstract, whereas the latter were discarded because the contradiction was too high.[31] Farm labourers are included as part of *manual* tasks. Armed forces occupations are discarded.

The mapping from occupations to tasks is presented in Table 3.3, for each standard. Instead of presenting the code in 4 digits for all occupations, whenever a whole group belongs to a given category, the group identifier is given. For example, if a 1 is shown it means all the occupations within that major group belong to that category. The correlation between the two SOC-based tasks is of 90% whereas between SOC and ISCO is around 77%. This difference might reflect the fact that the SOC – created by the ONS – captures better the UK occupations than the ISCO, which is designed to be a global standard. The final variable used for tasks categorisation is the one based on SOC00. Whenever this is missing, it is replaced with SOC90. If both are missing, the ISCO88 is then used.

[30] These are indicated at the bottom of Table 3.3.

[31] Many workers in this category were classified as manager in one standard (e.g. ISCO88), and producers in another (e.g. SOC). Expectedly, there is a blurred line between a farm producer and an owner; they are abstract as managers but routine as producers.

Table 3.3: Task-Occupation Mapping

Task	Occupations	ISCO88 (Four-digit)	SOC90 (Three-digit)	SOC00 (Four-digit)
Abstract	Managers, legislators, and senior officials	1	10-14 152-155 16-19	1 (except 1171)
	Professional occupations	2	2	2
	Associate professional and technical occupations	311-313 315 3222-3224 3226 33 3411-3413 3416-3419 345 347 348	30 32 34-39	31 32 34 352 353 355 356
Routine	Protective service occupations	516	610-613	3312-3319
	Administrative and secretarial occupations	4	4 691 94	4 921
	Sales and customer service occupations	3414 3415 52 911	7	354 621 7
	Skilled trades occupations	612 72-74	51-59	51 52 54 (except 5434)
	Process, plant and machine operatives	81 82 833 834	80-86 884-895 897 899	811-813 822
	Transport drivers and operatives	314 831 832	33 631 87 880-883	351 821
	Construction and extraction occupations	71 931	50 896 898 910 922-924	53 814 912
	Other routine occupations	342-344 932	911-919	913 9149
Manual	Personal service occupations	3460 511-515	614 615 619 62 630 66 67 690 699 950-954 956-959	5434 622-629 922 923
	Agricultural, fishery and related labourers	6112-6114 614 615 62 92	900-904	911
	Health service occupations	3221 3225 3227-3229 323 324	64 65	61
	Other manual and service occupations	912-916	920 921 929 93 954 990 999	9141 924 925
Discarded	Armed Forces	0	150 151 600 601	1171 3311
	Farming managers	1311 6111 6130 9211	160 169	5111

Note: Codes other than 4 digits means that the whole sub-group belongs to a given category. For example, 7 means all the occupations within that major group (7111, 7112, etc) belong to that category. The ISCO88 classification can be found here: html://www.ilo.org/public/english/bureau/stat/isco/isco88/major.htm. That website also includes the official description of each occupation and their respective tasks involved. The SOC90 and SOC00 classification can be found here: http://www.camsis.stir.ac.uk/occunits/distribution.html#UK. Official documentation describing each occupation is available at the ONS website.

Task descriptive statistics

Having explained the method used to map workers' occupations into tasks, it is possible to provide some basic descriptive statistics about the latter. The distribution of task in the sample is fairly uneven. Around 34% of all observations (i.e. worker-year pairs) correspond to workers in abstract-based occupations, whereas 49% belong to routine occupations. Just 17% account for manual tasks. This is in line with results in Holmes and Mayhew (2015) (Table 3.1). Figure 3.6 shows the evolution of the log real monthly wage per task group, over time. There is a clear differentiation between average wage among tasks, in a perhaps reasonable pattern. All tasks show a relative increase in wages over the period (indicated by the slope of the curves). Interestingly, abstract tasks show the lower growth rate, in contrast with the evidence of increasing inequality.[32] Figure 3.7 shows the frequency distribution of wages over all the period, compared with that of each task. In a sense, if the only determinant of wages were ability and the wage rate, this graph would offer the solution to the problem at hand. So, we can see these comparisons as the "unconditional" task distributions. They indicate that there is a fairly good distinction between tasks in terms of workers sorting. Manual workers tend to be at the bottom of the wage distribution, whereas abstract workers occupy the top. Again, these graphs are "unconditional" as they do not adjust for experience, education level, etc.

Figure 3.8 shows the proportional type of degree held by workers on each occupation. The sum between tasks and within degrees is 100%. University degrees are clearly related to abstract jobs. Nursing qualification are divided between abstract and manual. This is because nurses are classified as abstract, whereas assistant nurses and related are classified as manual (Cortes, 2016 follows a similar approach). No qualification is dominant in routine. This highlights the aforementioned overlap between routine and manual, which is also found in the estimations later.

[32] If the top 0.5% is kept, the increase is slightly larger, although not to reverse the difference with other tasks.

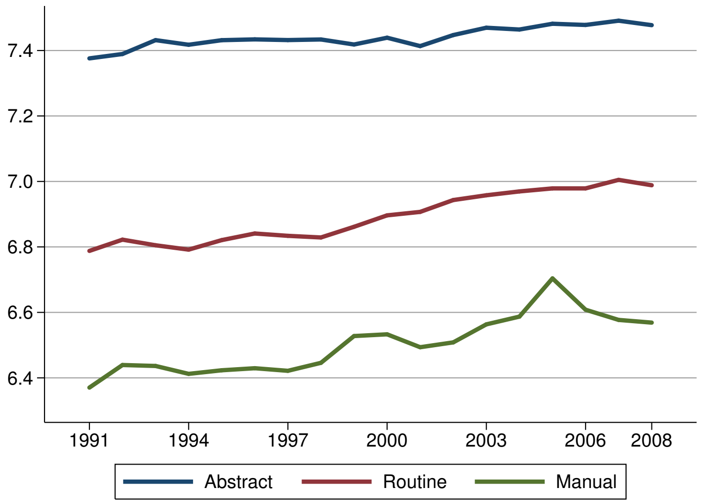

Figure 3.6: Evolution of average log real monthly wage, per task

Evolution of Tasks over time

Out of 15,843 individuals in the sample, 36% are observed in two tasks and 8% are in three tasks. Thus, the number of occupational switchers in the sample is 44%. This means a considerable portion of abilities will not be estimated using the Mincer approach. The imputation exercise required later is then of great importance.

Figure 3.9 shows the evolution of employment per occupational group as a proportion to total employment. Employment in routine occupations have been falling considerably, whereas abstract occupations have seen most of the increase. If anything, manual occupations have fallen during the period. This figure contrasts with the evidence of polarisation shown in Chapter 2. It tells us that the BHPS and LFS samples are not capturing the same proportion of the labour force. As the next section shows, the BHPS sample suffers from selection problems, whereas the LFS sample renews every five quarters, enhancing its population representativeness. The developments in Figure 3.9 are due to several factors like sample entry and exit, and in-sample switching. Moreover, these represent net effects, whereas gross mobility happens in all directions (see Table 3.4 later).

Transition between occupations is explored through a transition matrix, presented in Table 3.4. This table has to be read in rows. Thus, 87.1% of those who have ever been

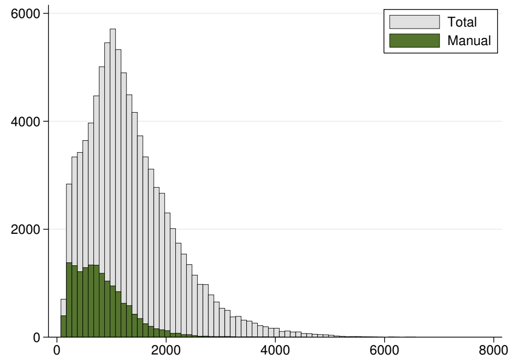

(a) Manual occupations

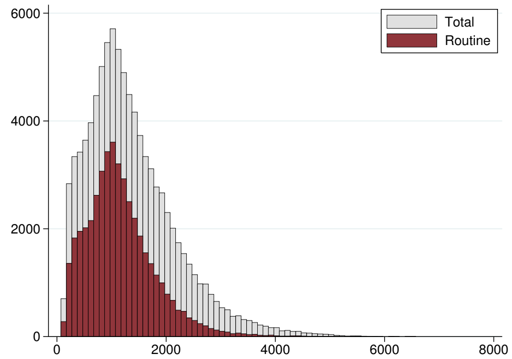

(b) Routine occupations

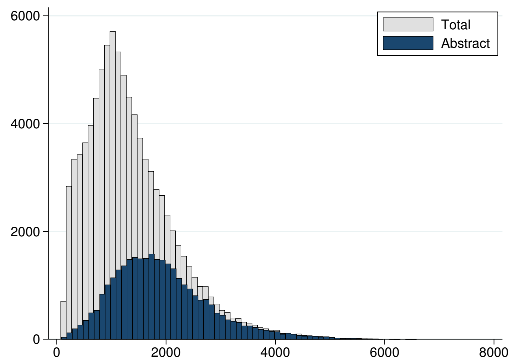

(c) Abstract occupations

Figure 3.7: Wage frequency distribution for each occupation versus the total wage frequency distribution. The green areas add up to the grey area.

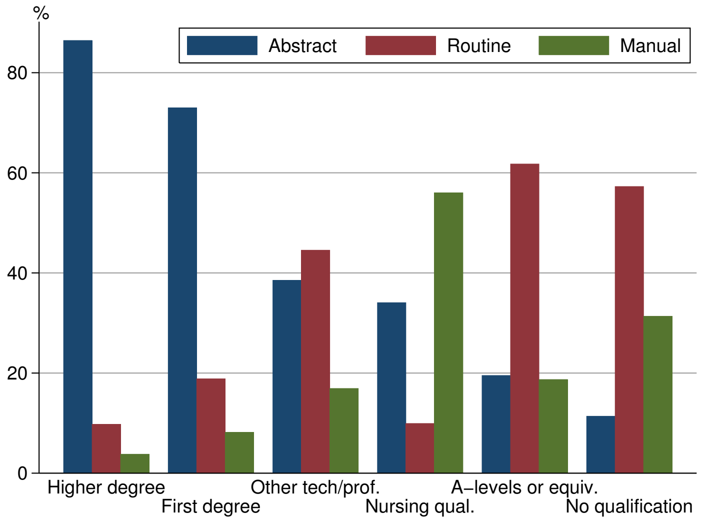

Figure 3.8: Highest qualification by task, average over period

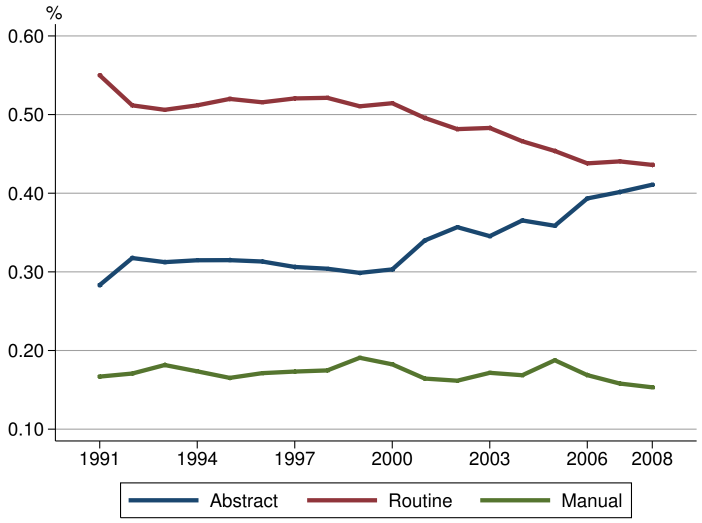

Figure 3.9: Share of total employment for each task

in the abstract occupation remained there in the next year, whereas 9.9% moved to a routine task in the next period, **on average across years**. It shows that downward mobility to manual occupations is quite limited compared with other transition types. Workers in routine occupations tend to move upwards. There is yet some switching from manual to abstract. Notice that this table does not indicate absolute levels of switching, but relative to the number of workers on each occupation.

In summary, the mapping between occupation and tasks, albeit simplified, seems to be a good enough approximation to the stylised model under study. A more refined differentiation between manual and routine might be beneficial. Yet, some ambiguity is

Table 3.4: Workers' transition between tasks in two consecutive years, average over period (%)

		To			
		Abstract	Routine	Manual	Total
From	Abstract	87.1	9.9	3.0	100
	Routine	8.7	87.9	3.4	100
	Manual	9.8	9.7	80.5	100

bound to remain when using a 1-to-1 mapping between tasks and occupations. This is a price to pay if one wants to identify ability (as said earlier, a more complex mapping does not allow such identification).

3.7 Estimation results

3.7.1 Estimation of coefficients

First, I present the results from each estimation method. All models produce equivalent signs, in the expected direction. Some coefficients are quite similar across models (e.g. full-time, for both gender), whereas other are quite different. As expected coefficient from FE and HT estimations are quite similar, and in most of the cases statistically indistinguishable (based on simple confidence interval overlap analysis). Coefficients from FE-2SLS are generally different, and closer to zero. Moreover, their standard errors are larger, making some of them less or no significant at all. These two problems (i.e. downward bias and large standard errors) are however common in 2SLS estimation (e.g. Wooldridge, 2010, §5.2.6). They might partly explain the insignificance of education in FE-2SLS, a sharp contrast with other methods. Moreover, as reported in the table, the null hypothesis of exogeneity of education is not rejected. This could be because education is in fact not endogenous, or due to poor instruments. The latter seems not to be the case though, as both under and weak identification tests are rejected at 1% in favour of the instruments used. It is also of interest the apparent over-estimation of education in RE compared with FE and HT. This might reflect the omission of ability, likely to be correlated with education, providing a positive bias. A standard Hausman supports this correlation by favouring FE over RE, as the table indicates.

Table 3.5: Estimation for log of monthly pay, 1991-2008

	Model			
Variable	**RE**	**FE**	**Hausman-Taylor**	**FE-2SLS**
White	0.0748***		0.0684***	
	(0.017)		(0.015)	
Female	-0.399***		-0.412***	
	(0.026)		(0.037)	
Experience	0.102***	0.123***	0.112***	0.0978***
	(0.003)	(0.005)	(0.004)	(0.009)
Experience × female	-0.0252***	-0.0281***	-0.0292***	-0.0258**
	(0.005)	(0.008)	(0.007)	(0.012)
Experience^2	-0.538***	-0.647***	-0.612***	-0.438***
	(0.028)	(0.035)	(0.035)	(0.059)
Experience^2 × female	0.0905*	0.161**	0.135**	0.157
	(0.046)	(0.064)	(0.062)	(0.096)
Education	0.0312***	0.0136***	0.0141***	-0.0032
	(0.002)	(0.003)	(0.003)	(0.066)
Education × female	0.0137***	0.0001	0.009**	
	(0.002)	(0.004)	(0.004)	
Full-Time	0.521***	0.497***	0.508***	0.486***
	(0.020)	(0.024)	(0.025)	(0.033)
Full-Time × female	-0.375***	-0.359***	-0.368***	-0.368***
	(0.022)	(0.026)	(0.029)	(0.035)
Hours per week	0.0129***	0.0119***	0.0122***	0.0088***
	(0.001)	(0.001)	(0.001)	(0.001)
Hours per week × female	0.0182***	0.0178***	0.0180***	0.0219***
	(0.001)	(0.001)	(0.001)	(0.001)
Public Sector	-0.0091	-0.0086	-0.0075	-0.00910
	(0.008)	(0.012)	(0.009)	(0.014)
Public Sector × female	0.0800***	0.0523***	0.0621***	0.0391**
	(0.011)	(0.014)	(0.012)	(0.019)
Trade Union	0.0510***	0.0420***	0.0452***	0.0433***
	(0.005)	(0.007)	(0.007)	(0.010)
Trade Union × female	0.0165**	0.0038	0.0084	0.0003
	(0.008)	(0.009)	(0.009)	(0.014)
Tests: (p-values in parenthesis)				
Hausman (χ^2)	584.2***			
	(0.000)			
Exogeneity (F)				0.241
(Davidson-MacKinnon)				(0.62)
N	90,780	90,780	90,780	55,268

* p <0.1, ** p <0.05, *** p <0.01

Notes: standard errors in parentheses. All models include marital status, firm size, industry and region controls. The 2SLS model uses education of partner and father as instruments. The standard errors for HT are bootstrapped with 100 replications.

Regarding the estimated coefficients, these are not trivial to compare with other studies because their levels depend on the exact definition of independent and dependent variables, and also on the base group categories used. The effect of education might be of particular interest though. For the HT estimates, an extra year of education is associated with an average increase on real wages of around 1.4% for male, and 2.3% for women. In the case of FE, the increase is of 1.4% for both gender. This might be smaller than other findings. For example, the mentioned twin-based study for the UK by Bonjour et al. (2003) find that an extra year of education increased average **nominal hourly wages** by 7.9% (IV estimates, Table 2). Their standard error is however much greater, leading to a 95% confidence interval of (1%, 15%). Accounting for inflation does bring them closer (it was around 2.3% on average over our sample period). Still, there are considerable differences between the two approaches, including a complete different sample. Perhaps an important difference is that they measure education based on an assumed duration of degrees. Here however, education is calculated from self-reported age at which education finished (adjusting for mature students). In any case, this is just one example of how hard it is to compare results across studies, reason why a thorough comparative analysis is not carried out here.

As highlighted in Section 3.4, for FE models we can run a second stage equation to identify the coefficients of time-invariant observables. In this case, it is race and gender. As mentioned, the quality of this estimation depends heavily on the consistency of the estimation of the fixed effects, which depends directly on the length of the panel. The estimates are shown in Table 3.6, where reported standard errors are bootstrapped with 100 repetitions. For contrast, they include the estimation using RE and HT methods, from Table 3.5. Furthermore, FE estimates include a model with and without constant, for reference. The latter is the one corresponding to the theoretical model outlines in Section 3.4.

The results indicate considerable differences between estimates, particularly for race. On the one hand, RE and HT provide statistically indistinguishable results. Since HT accounts for potential correlation between education and unobserved effects, this might be an indication that race is not correlated with unobservables. Otherwise, RE estimates would present evidence of bias. On the other hand, FE and FE-2SLS estimates for race

Table 3.6: Estimation of time-invariant regressors, 1991-2008

	RE	**FE**		**Hausman-Taylor**	**FE-2SLS**	
Variable		No constant	Constant		No constant	Constant
Constant			0.207***			0.349***
			(0.032)			(0.039)
White	0.0748***	0.2119***	0.0113	0.0684***	0.3108***	-0.0403
	(0.017)	(0.005)	(0.033)	(0.015)	(0.006)	(0.039)
Female	-0.399***	-0.4695***	-0.4816***	-0.412***	-0.7067***	-0.7335***
	(0.026)	(0.009)	(0.01)	(0.037)	(0.008)	(0.013)
Adj-R^2	0.72	0.24	0.24	–	0.37	0.38
N	90,780	22,232	22,232	90,780	12,039	12,039

* p <0.1, ** p <0.05, *** p <0.01

Notes: Standard errors in parentheses. The standard errors for FE, HT, and FE-2SLS are bootstrapped with 100 replications.

without a constant are considerably higher than the HT and RE models. In the case with a constant, race becomes insignificant. However, the constant is highly significant, and the addition of both yields almost the same coefficient than the model without constant. This is likely an outcome of the high dominance of whites in the sample (97% of all individuals), as Table 3.1 indicates.[33] It is also worth noticing that the number of observations used in FE-type of models is considerably smaller. This could be behind the observed disparities in results. As argued above, race is likely to be uncorrelated with unobservables, so there assumption required for the second stage identification of time-invariant variables under FE-type of models might not be violated. All in all, evidence from HT estimation indicates an approximate 7% higher real monthly wage for whites compared with non-whites, *ceteris paribus*. The magnitude itself is hard to compare with other studies (due to different definition of variables), but in line with evidence available for the UK (e.g. Brynin and Güveli, 2012 and Dustmann and Theodoropoulos, 2010).

Results are quite more homogeneous for the gender dummy (except under FE-2SLS). There is no significant difference in models with and without constant, due to the even

[33] Bootstrapping methods allow to estimate the potential bias in the point estimates obtained from the original sample (those shown in the table). Yet, for race the bias is less than 2% of the point estimates, meaning they alone cannot explain the differences.

representation of gender in the sample. On average, females receive lower wages than men.[34]

To conclude, unobserved heterogeneity exists and it is correlated with other endogenous variables. As such, estimations from the RE models are inconsistent. FE and HT perform fairly similarly, except for time-invariant factors. The latter model might be more trustworthy on this respect. The FE-2SLS showed no evidence of the endogeneity of education with respect to the idiosyncratic error, and so external instruments are not needed. Due to intrinsic biases, this method is inconvenient for the estimation of ability, which relies on unbiased and consistent estimators. Hence, in what follows the emphasis is put on both FE and HT estimations.

Wage rates

Panel (a) in Figure 3.10 shows the evolution of the estimated wage rate per occupation over time, in logarithmic terms, for the HT estimation.[35] Notice the difference in levels, as expected. All wage rates increased over the period. Although not shown, a 95% confidence interval indicates manual and abstract occupation rates of change are essentially the same. This is not the case for routine occupations. The growth is slower, and statistically nil after 2002. The magnitude of these changes is also relevant, as compared with the evolution in observed wages per task (shown earlier in Figure 3.6). For example, the increase in manual occupation wages is of the same magnitude as the increase in its wage rate. This does not imply other factors did not change over the period, but it does mean they more or less cancel out. Panel (b) shows the relative change between abstract and manual log wage rates relative to routine tasks. There is a consistent increase, albeit mild over the period. These results are particularly interesting in light of the analysis carried out in Chapter 4. Namely, the increase in relative productivity of routine tasks with respect to other tasks – due to technological change, which together with relatively low elasticity of substitution between tasks leads to an increase in the wage rate of manual and abstract tasks relative to routine tasks.

[34] More precisely, they receive 40% lower if thy belong to the base group. Yet, this is an unlikely group including women without experience nor education, working part-time, in the private sector, without trade union, single, and with zero hours of work.

[35] The estimation of the wage rate is quite similar in both FE and HT estimations. More precisely, their correlation is 99%. The grand difference is that HT enables us to compare their levels, since these rates are estimated for a common task year (in our case, manual task for 1991).

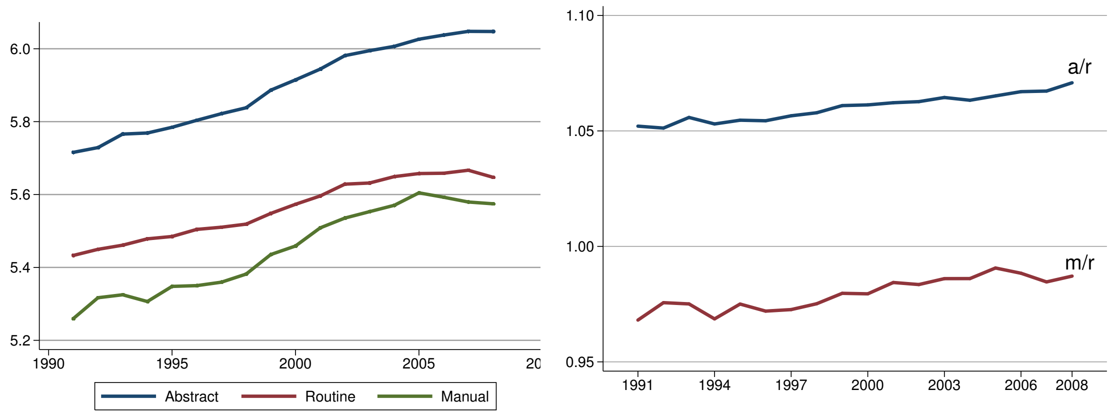

(a) Log wage rates for each task

(b) Abstract and manual tasks' log wage rates relative to routine task's log wage rate

Figure 3.10: Evolution of estimated log wage rate effect per occupation, HT estimation

3.7.2 Unobserved effects

As mentioned in subsection 3.3, using the estimated coefficients for all regressors we can generate a prediction of the unobserved effects, c_I. Strictly speaking, this component encapsulates all time-invariant unobserved factors affecting a worker's wage, specific to a given occupation. If there are no other time-invariant observables omitted from the model, and ability is indeed the only unobservable factor affecting wages, we can conclude that the predicted errors correspond to our destination variable, ability. Whether we can actually identify ability or not, that is given by the statistical properties of the estimators (i.e bias and consistency).

The correlation between predicted effects for FE and HT estimations is 80%. Since ability prediction depends on the estimated coefficients, this not-so-high correlation might be reflecting the disparities already highlighted in Table 3.6. Their contrast is more evident in Figure 3.11, which depicts the evolution of **average** of predicted log ability, for each task, for both FE and HT estimations.[36] These averages are calculated **only** for those who are currently employed in the respective task in a given period.[37]

[36]It is important to notice that, since each estimation includes a constant, these estimated effects have been demeaned. This is not a problem *per se*, because subtracting a constant from a distribution only changes its location but not it's shape or scale. When making comparisons among workers, the interest is also on *relative* abilities.

[37]Recall that ability is calculated for all occupations at which workers are observed throughout the sample.

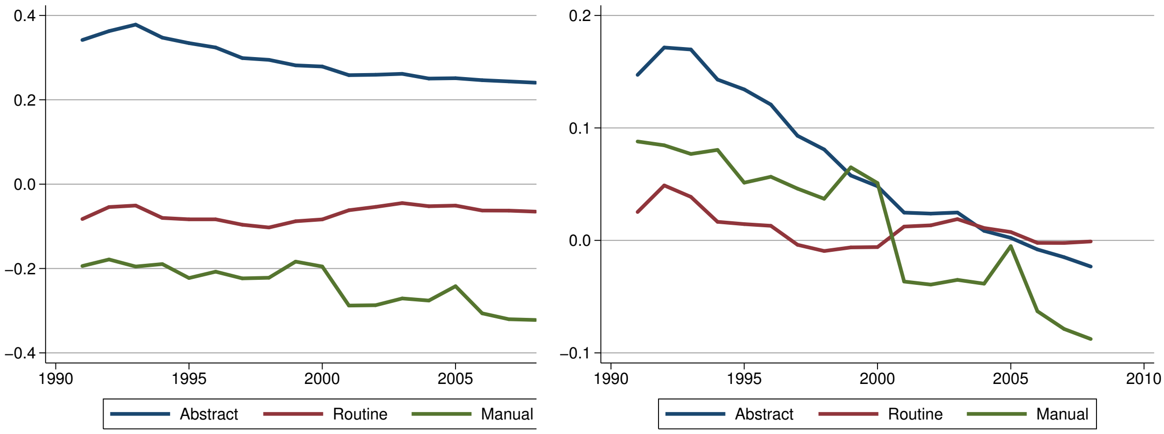

(a) FE estimation (b) HT estimation

Figure 3.11: Evolution of predicted mean log ability per occupation

Since predicted workers ability is constant over time, fluctuations in average log ability observed in Figure 3.11 are necessarily due to (1) entry and exit from the sample, and (2) switching in and out of occupations. The former is expected to be important, in light of the evidence shown in Panel (b) of Figure 3.5. For example, if workers with higher than average ability are leaving the sample, mean ability falls. Similarly, if lower than average ability workers are entering to a given occupation, average ability also falls in that occupation. The effect of switching is more complex, and depends on the extent of absolute advantage in the model (recall **Result 3**). Switching can alter average tasks' ability because those leaving (entering) a task might have systematic ability differences with respect to stayers (incumbents).

The analysis of these factors are presented in Figure 3.12. They show the average ability of those who enter into and leave the sample from a given task, plus the mean ability of those who switch into or from that task, in comparison with those who are already in that given task (the horizontal, black line). Notice these graphs do not capture the magnitude of each change. For example, the number of those switching might be much greater than those entering. Still, together with Figure 3.11 they can reveal the overall source of ability changes. Panel (a) on Figure 3.12 shows that the fall in mean abstract ability is mainly due to those entering the task from outside and inside the sample having lower ability than current workers. Panel (b) indicates that entrant and leavers play contradicting forces in the overall trend in routine ability. Switchers into

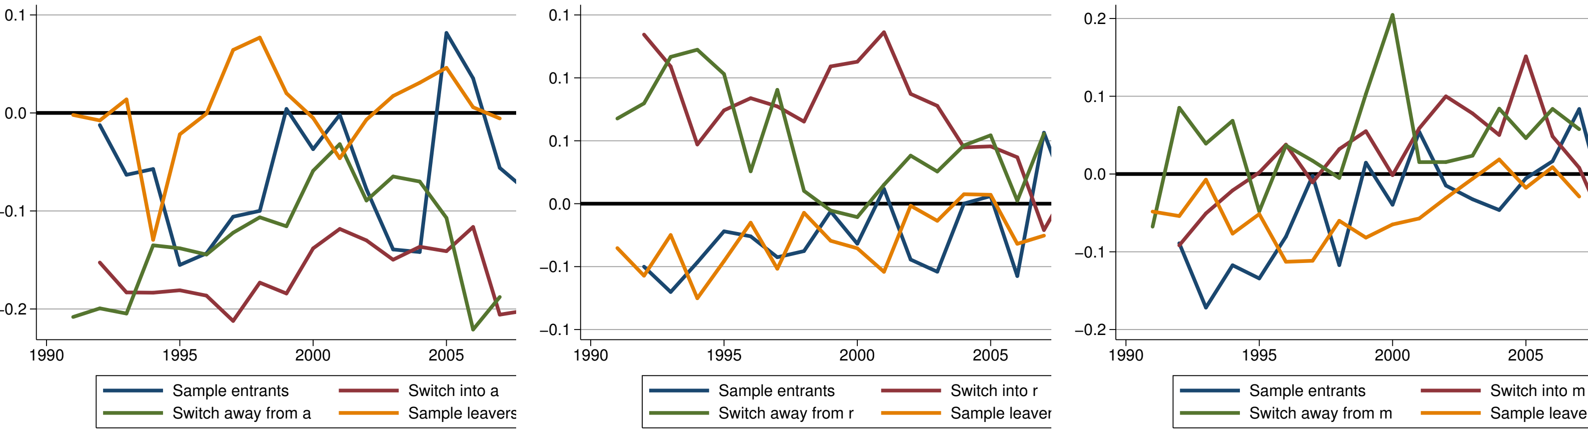

(a) Abstract task, HT estimation

(b) Routine task, HT estimation

(c) Manual task, HT estimation

Figure 3.12: Relative ability of each subgroup compared with those who are currently employed on a given occupation

and away from routine occupations also affect overall ability in opposite ways. Given that the mean of routine ability is relatively stable, these forces more or less balance in the aggregate. Finally, given the overall fall in average ability in manual occupations, Panel (c) indicates that switching into this task and those entering from outside might be the largest drivers of this trend. Overall, Figure 3.12 highlights the multidimensional phenomena behind the changes in occupations' ability across the sample over time. The contribution of entry and exit is significant, just as that of switching. The particular effect of them varies across tasks.

3.7.3 Workers' ability in observed occupations

The distribution of observed ability for each task is shown in Figure 3.13, for the year 2001 (when sample is the largest), and using HT estimates. These are histograms, or relative frequency graphs, where the area add up to 1. Notice these distributions do not include the ability of those who are not observed in a given occupation. As such, they do not represent our final destination. They are related to Figure 3.17 in Section 2.3, and to the examples shown from other literature, in Section 3.2. They are all positively skewed, perhaps resembling a log-normal distribution, in contrast with Cortes (2016), who's distributions were if anything negatively skewed, and quite different from

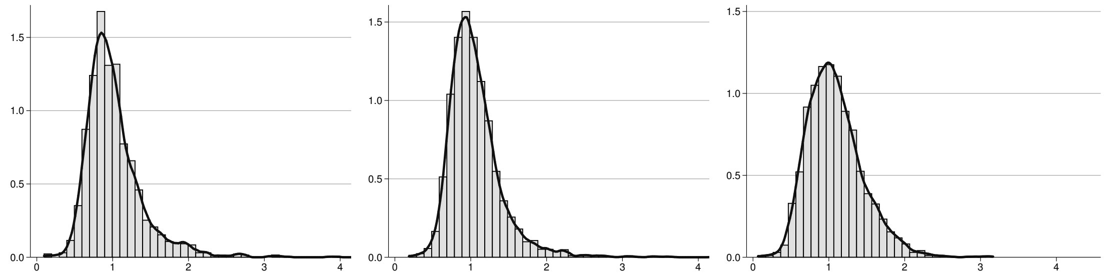

(a) Manual task (b) Routine task (c) Abstract task

Figure 3.13: Ability distribution for workers employed at the respective occupation, for year 2001, HT estimation

log-normal. Notice that these distributions are approximately centred at 1 because predicted **log** ability has mean zero (and $e^0 = 1$). What is informative then is not their level but shape and dispersion. Importantly, notice they cover a relatively wide spectrum of ability. This is, for each of them there are very low skilled individuals, just as there are also very high skilled ones, relative to the rest of the workers. This contrasts with the stylised case of Figure 3.17 (absolute advantage), where distributions were clearly separated. This is relevant later.

3.7.4 Further analysis of observed ability

Uncertainty of individual ability prediction

Section 3.3 left clear that the prediction of ability c_I is only consistent when $T_I \to \infty$ (albeit still might be unbiased). Since the panel is unbalanced, and the number of years switchers spend in different occupations is heterogeneous, ability is predicted with varied degrees of uncertainty. It is of interest to quantify this uncertainty, and relate it to the number of years used to predict ability. Also, we can compare uncertainty in both estimation methods. This is accomplished by bootstrapping the model estimation, and recording the prediction of ability for each repetition, for each cluster unit I. The standard deviation of these many predictions is then an indication of the degree of uncertainty with which individual ability is calculated. The results of this exercise is presented in Figure 3.14, for a bootstrapping with 1,000 repetitions. It shows the average of the standard deviation across individuals that are in an occupation a given number of years, for all possible years – from $T_I = 1$ to $T_I = 18$, and for both models.

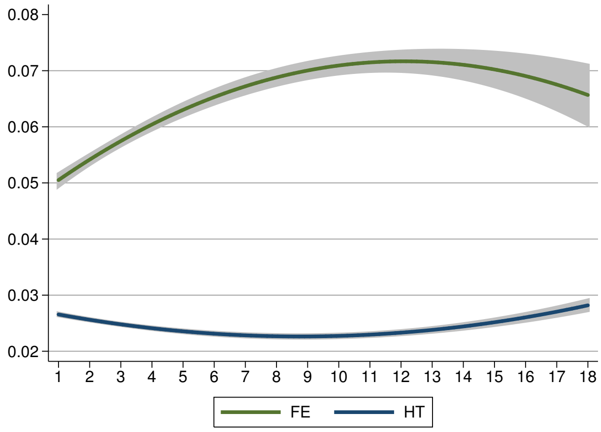

Figure 3.14: Average standard deviation of prediction of log ability across individuals for levels of T_I, FE and HT estimations

A 95% confidence interval is also included. It is clear that the level of uncertainty in FE-based predictions is much higher than with HT – for some years, standard deviation is three times higher. Uncertainty for different years within models is also greater for FE model – between 0.05 and 0.07, compared with HT estimates – between 0.02 and 0.03. Interestingly, the pattern across years is the inverse. The lowest uncertainty in predictions from HT estimations is for $T_I = 9$, whereas the largest uncertainty in FE-based estimates is for $T_I = 12$. It was expected for variability of estimates to fall as T_I increases – as the predictor's variance in equation (3.6) indicates. This is clearly not the case for FE estimates. However, it does hold for HT from $T_I = 1$ until $T_I = 9$ – with later increases in uncertainty being rather low. This is another argument in favour of HT over FE estimates.

Is ability correlated with education and wages?

The objective of using fixed effects is to eliminate unobserved heterogeneity that might be correlated with some observables included in the model. If ability is in fact behind this heterogeneity, we might expect it to be correlated with education and, perhaps, with wages (after controlling for education). This turns out to be then an informative test about the nature behind the estimated unobserved effects.

The first row in Table 3.7 shows pairwise correlations between education and the **level**

Table 3.7: Pairwise correlation between ability and selected variables, HT estimation

	Ability		
Variable	**Manual**	**Routine**	**Abstract**
Education	0.28	0.15	0.17
Monthly wage	0.57	0.61	0.61
Father's education	0.14	0.12	0.16

of estimated ability, for each task for the HT estimation.[38] At first, results indicate a fairly low level of correlation between ability and the education regression, ranging from 15 to 28%. Moreover, it might seem puzzling that correlation is the highest for workers in manual tasks. Nevertheless, at least since Becker (1964) seminar contribution, economists have acknowledge the fact that educational decisions not only depend on ability but also on access to financial resources – including own wealth. This is, those who pursue higher education and further studies are not always the most qualified, but those who have better access to networks, information and funding.[39] In other words, ability is a crucial determinant for the highest educational level someone in manual occupations can achieve compared with those in higher-paying occupations. This result is then consistent with observed lack of equality of educational opportunities among people with similar ability (see footnote at bottom of this page).

Nevertheless, several arguments go against a very strong correlation between education and ability in the first place. On the one hand, the offer of further education is very heterogeneous with respect to quality. This means that people of quite different ability levels can pursue the same degree in institutions of different quality. Additionally, different degrees might have the same length (e.g. adult nursing and geography, both

[38] In the model ability enters in log terms. Thus, in terms of potential endogeneity of ability due to correlation with education, its log definition might be of more relevance. Although not shown, correlations in logs are roughly the same as those presented in the table, if not higher. Regarding the FE model, correlations are around 5 percentage points higher for education, 15 percent points lower for wage, and roughly equivalent for father's education.

[39] For example, Lovenheim (2011) shows how an increase in household wealth due to home price appreciation positively influences the likelihood of college enrolment, the quality of such college, and the probability of graduating. Oreopoulos and Dunn (2013) use a randomize field experiment to show that access to information about post-secondary education among high-school students increases their willingness to pursue further education, their expectations about own future achievement, and knowledge that access to financial aid was possible for them. Most of the literature on intergenerational mobility highlight these aspects too. For instance, see Mazumder (2005) and Hanushek, Leung and Yilmaz (2014). For a survey on the topic, see Bowles and Gintis (2002).

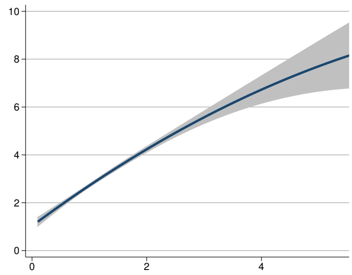

(a) Manual ability, quadratic fit

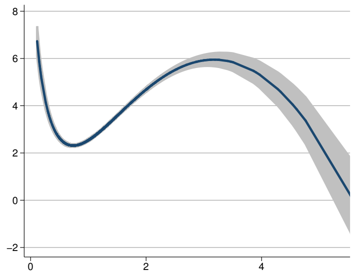

(b) Manual ability, cubic fit

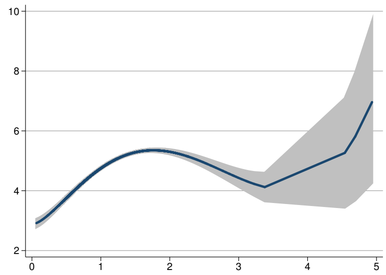

(c) Abstract ability, cubic fit

Figure 3.15: Non-linear relationship between education (y-axis) and ability (x-axis), HT estimation (with 95% confidence interval)

lasting three years). In consequence, using years of education will not capture qualitative heterogeneity within degrees or different type of degrees. On the other hand, there are upper limits to the number of years of studies very clever individuals can achieve. Besides, not every highly intelligent student is expected to pursue highest educational degrees like PhD or a Post-Doctorate. If true, these issues lower the correlation between education and ability.

Figure 3.15 explores some potential non-linearities. This is by no means a rigorous analysis, but it illustrates the point. For example, a quadratic fit for the relationship between level of ability and years of education seems quite linear. However, a cubic fit indicates that is likely not the case. Similarly, a cubit fit for abstract ability indicates possible non-linearities. The truth is that the fit vary considerably according to the degree of polynomial used and the inclusion or not of "outliers". Yet, what emerges from these is that the relationship between education and predicted ability is unlikely to be linear. In consequence the low correlation found in Table 3.7 should not be a surprise. Finally, the pairwise correlations and polynomial regressions shown above are not accounting for any other factor that might play a role in the link between ability and education. In order to fully asses this relation, a proper study is needed.

Table 3.7 also presents the correlation between ability and wages – both in levels. While this could be capturing the role of education, or other factors related to ability, there is a much higher correlation – around 60%. This might be an indication that productivity at work is not solely a function of formal qualifications but of actual ability. Finally, own ability seems to be positively but rather weakly correlated with both father's education.[40] Although we might have expected a higher correlation, the same concerns with education apply here. The good news is that our predicted effect might not be capturing other unobservables like family background.

Does ability vary across occupations?

A central element of the theoretical and econometric model is the assumption of occupation specific ability. The analysis so far already shed light on clear differences on ability across occupations – in particular Figure 3.11 and Figure 3.12. A more formal test was outlined in Section 3.4. Namely, a Hausman test that evaluates the validity of the clustering by comparing estimation results under both standard and occupation-spell clusters. Implementing the test upon both FE and HT estimations overwhelmingly rejects the null hypothesis that ability does not vary over occupations (p-value of 0.000).

3.8 Estimating unobserved ability

The previous section produced an estimate of "observed" ability for each task, for every year. Our ultimate goal is to produce an estimation of the complete ability distribution, which require to estimate ability in occupations in which workers were never observed. The only way to achieve this is to use a model which tells us some information about what such unobserved ability might be. This section presents such model, and shows how unobserved ability can be estimated from it. The most important take away from this part is that this exercise is in itself embedded in deep uncertainty; we cannot know in exactitude what unobserved ability might be. The best we can achieve is to produce some bounds and match some aggregate correlation levels.

After this theoretical analysis is carried out, a empirical evaluation of the key assumptions of the model is conducted, using the estimation results from the previous section. The goal is to justify the usage of such model to estimate unobserved ability. Finally, once this is done, unobserved ability is estimated, concluding the exercise.

[40] Father's education is chosen due to its relatively low number of missing observations compared with mother's education. When the former is not available, the latter is used instead.

3.8.1 The model

Chapter 2 presented a general equilibrium model where workers sort into occupations based on comparative advantage. The key markets in the model are the market for relative tasks. The demand for tasks comes from firms, whereas the supply of tasks comes from workers' ability. Their equilibria defines final employment shares on each occupation and tasks' wage rates, ω_j. The wage a worker is paid depends on the wage rate of the selected occupation and the worker's ability in that occupation (see equation (2.26) in the previous chapter).

In order to produce an econometric model of wages, we need to treat wage rates as exogenous to workers' decisions.[41] The model above assumes this, as the allocation decision by an individual worker has no effect on ω_js. This is also an empirically reasonable assumption. In a large economy as the UK, occupational choices by individuals should have no effect on market wage rates. Following this assumption, this section presents a partial equilibrium perspective of the general equilibrium model in the previous chapter. In particular, wage rates are exogenous; any change in demand are expected to be reflected in those wage rates. This is not just a simple repetition of previous chapter's model. Here, particular emphasis is put on the model's observational consequences for an econometrician. Building upon this model, the next section derives the econometric equation to be used in the estimations. Naturally, the production side of the model is neglected, because it has not relevance for the labour market equilibrium and wage determination. More precisely, once workers' allocation is decided, production and consumption follows through, without any feedback on the sorting mechanism.

Ability

Consider an economy populated by a continuum of workers of size $L = 1$. There are three tasks, or occupations, where workers can be employed. These tasks are: a *manual* task (m), a *routine* task (r), and an *abstract* task (a). Each worker i is endowed with a certain ability level for each task j, ability that is time-invariant and denoted η_{ij}. In

[41] If not, the econometric model might suffer from endogeneity, as some of the regressors – wage rates – are also an outcome of workers' choices. For more details, see Section 3.3.

other words, the ability set of worker i is $\{\eta_{im}, \eta_{ir}, \eta_{ia}\}$. Given worker's ability sets, there is a **distribution of ability** for each task among the population. An example of three log-normal distributions is shown in Figure 3.16, with ability level on the horizontal axis and the probability density function on the vertical axis.

(a) Manual task (b) Routine task (c) Abstract task

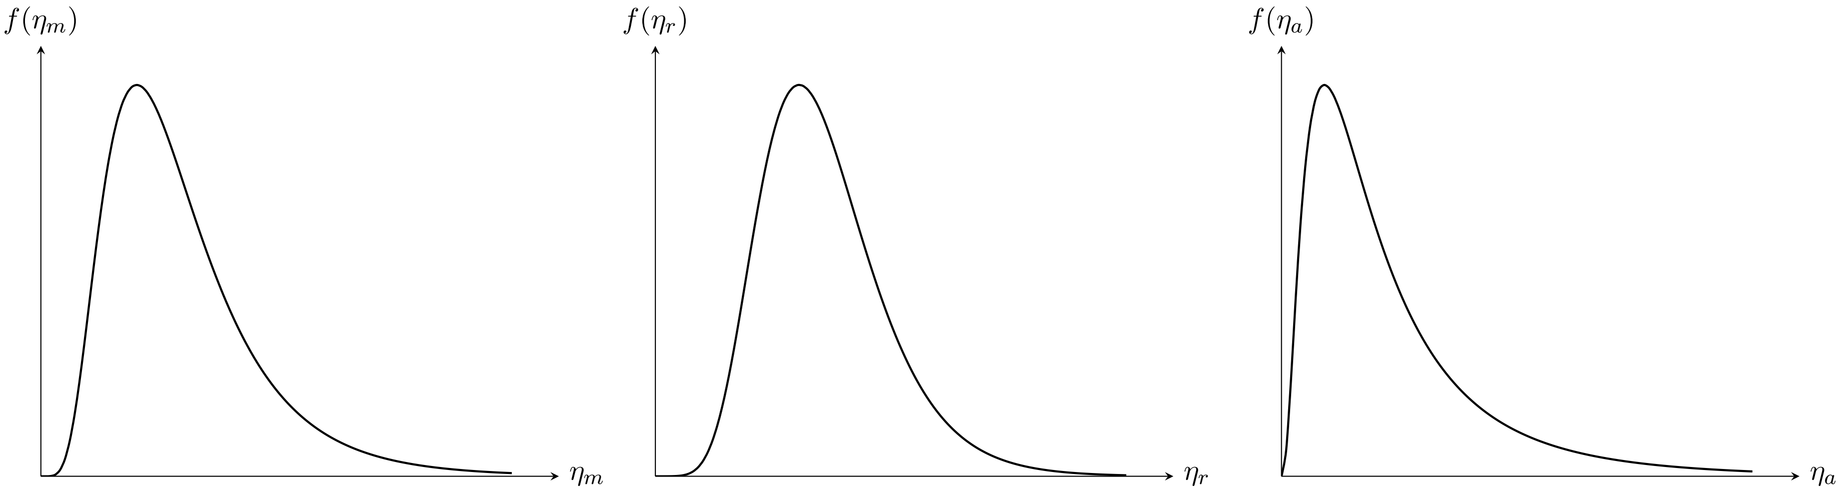

Figure 3.16: Distribution of workers' ability for each task

(a) Manual task (b) Routine task (c) Abstract task

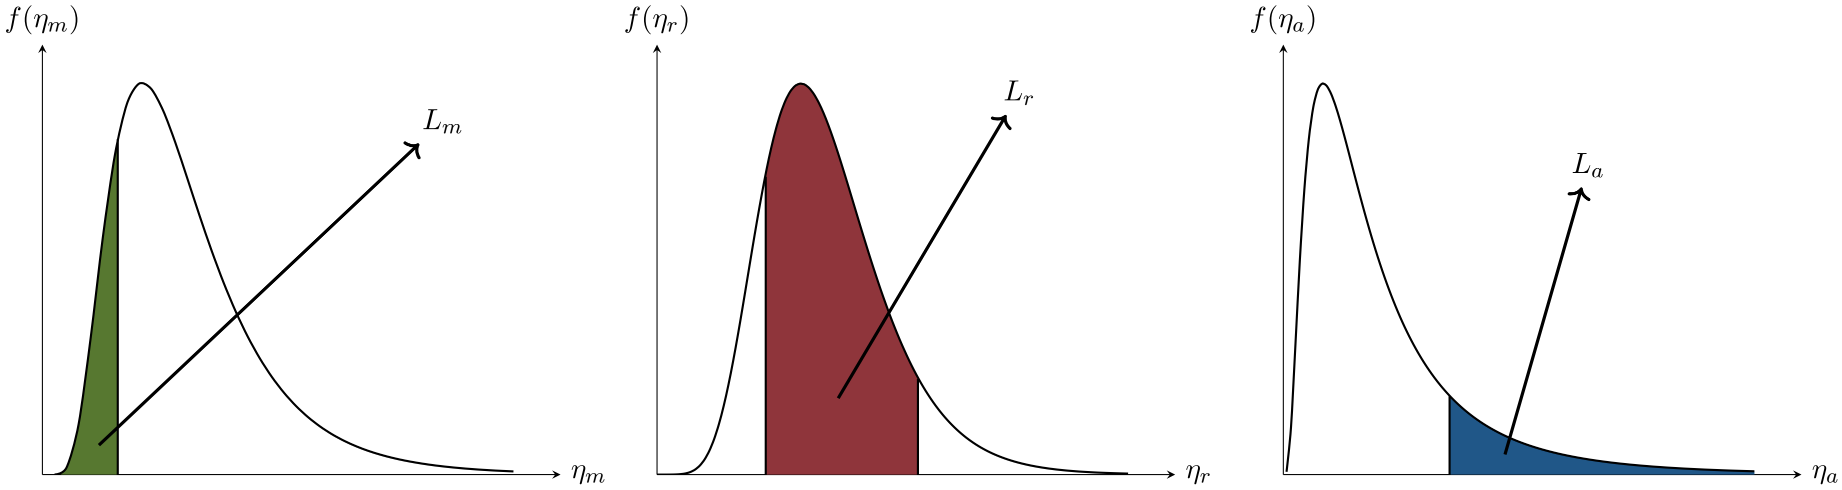

Figure 3.17: Example labour market allocation

Labour market equilibrium

Consider the following set of assumptions:

Assumption 0.

(i) The labour market is competitive. This translates into workers being paid their marginal product.

(ii) Workers utility only comes from wages.

(iii) There is perfect information. This is, workers know their ability and the wage they would get on each occupation.

(iv) There is full, costless mobility between tasks.

These have important consequences. **Assumption 0, i** leads to:

$$w_{ij} = \omega_j \eta_{ij} \tag{3.14}$$

where w_{ij} is the wage a worker i would receive if employed in occupation j – in real monetary terms, ω_j is the "wage per ability unit" in occupation j – also denoted wage rate, equal for all workers, and η_{ij} is worker's ability in occupation j. The rest of these assumptions imply that workers sort into the occupation that maximises their wage. This sorting leads to an equilibrium labour market allocation.[42] What would this allocation look like in Figure 3.14? Without further information, it is impossible to tell. To move forward, assume the following:

Assumption 1. *Absolute Advantage (henceforth AA): if a worker has higher ability than another worker for task j, he also has higher ability for all other tasks.*

[42] Again, this is a partial equilibrium setting. Chapter X studies the general equilibrium case, where the wage rate depends endogenously of both supply of and demand for labour.

In other words, if we order workers from low to high ability in the manual task, that order would also hold for routine and abstract tasks. AA is sometimes referred in the literature as *hierarchical* ability. **Assumption 1** allows us to use a simplified notation. Namely, we can order workers such that $i = 0$ is the lowest skilful worker and $i = 1$ is the most skilful worker. Because of AA, this order holds for each task. The next assumption build upon this ordering.

Assumption 2. *Manual-routine-abstract ordering: consider an ordering of workers from low to high ability, over the i dimension. The ratio of abstract ability over routine ability $\frac{\eta_{ia}}{\eta_{ir}}$, and the ratio of routine over manual ability $\frac{\eta_{ir}}{\eta_{im}}$ are increasing on i.*

This second assumption imply that the most (least) skilful workers for the abstract (manual) ability are employed in the abstract (manual) task, whereas those employed in the routine task are those of mid-range routine ability.[43]

Given these two assumptions, the labour market equilibrium is characterised as

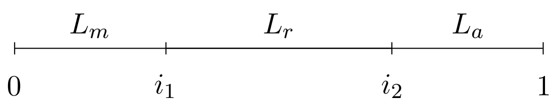

where i_1 and i_2 are the two workers who are indifferent between "adjacent" tasks. This equilibrium allocation can also be seen in Figure 3.17. Each distribution is partitioned in two sections, namely, those who are in that respective task and those who are employed elsewhere. The sum of the three coloured areas is one. Naturally, there is a close relation between the "cut-off" points on these distributions, given by equation (3.14). More precisely, they reflect the level of ability of indifferent workers. Notice that these graphs are equivalent of those in Figure 2.7, in Chapter 2, albeit from different angles. The latter show the labour market allocation using distributions' *cdf*, whereas here the equilibrium is shown from the distributions' *pdf* perspective.

Observed and unobserved ability

Imagine a representative, longitudinal sample is drawn from this population. Normally, ability is not observable. Yet, as described in the next section, it can be estimated

[43] See Chapter X for the proof this implication.

Table 3.8: Example of dataset with estimated ability from observed occupations

Worker	Task		
	Manual	Routine	Abstract
1	1.8		
2		3.2	4.0
3	2.0	2.3	
4		2.9	3.6
5			5.7
6	3.8	4.3	6.1
⋮	⋮	⋮	⋮

for occupations in which workers are observed. This is, an estimation of the bold areas in Figure 3.17 can be deduced from the data. For those who never switch occupations in the sample, ability is estimated for only one task. For those who switch during the sample period, ability is estimated for at least two tasks.

Since not everyone is a switcher, and those who switch rarely are employed in all three tasks, there is considerable missing information regarding workers' ability. An example of what this dataset might look like is given in Table 3.8, for six workers, where empty cells mean the worker has not been employed in that occupation. If our objective is to fully characterise these distributions, how can "observed" ability be used to impute "unobserved" ability?

First, notice that **Assumption 0** implies the following:

Result 1. *A worker is observed in the occupation which it is optimal for her/him. By implication, a worker is not in other occupations because it would be suboptimal. More precisely:*

$$\eta^u_{i,-j} < \eta_{ij}\frac{\omega_j}{\omega_{-j}} \tag{3.15}$$

where η^u_{i-j} is the ability of worker i in all unobserved occupations $-j$.

From the estimation of the model we can derive values for wage rates and for "observed" ability. Then, using equation (3.15), we can produce bounds for unobserved ability, for every task at which a worker is not observed, and **for every period**. Since ability is assumed to be time-invariant, the lowest boundary among the period is taken as the effective bound. Since ability is by assumption non-negative, unobserved ability for each worker is constrained between zero and the aforementioned bounds.

Now, at first, it is reasonable to expect that a worker's ability for a task he is not employed in might be lower than the ability of those who are already employed there. Otherwise, this worker could have switched. This would enable us to provide further constraints to workers' unobserved ability. Yet, this is not the case. Recall that sorting is based on comparative advantage, and not on absolute advantage.[44] As such, there is no particular constraint that can be imposed, unless we impose further assumptions. Thus, given the analysis so far, the only certainty is that given by **Result 1**.

Assumption 1 and **Assumption 2** lead to the following result:

Result 2. *The ability of all those who are never in the abstract occupation must be lower than the ability of those who are observed in abstract tasks. Similarly, the ability of all those who are never in the manual occupation must be higher than the ability of those who are observed in manual tasks. Finally, for workers who are employed in abstract (manual) tasks but never on routine ones, their ability in routine occupation must be higher (lower) than the ability of those who are observed in routine tasks.*[45]

This result can be seen in Figure 3.17. In particular, notice that the best for manual and routine tasks are not employed in those occupations, because it happens to be that those are also the best for abstract tasks. By virtue of **Assumption 2**, the best among all the population are relatively more profitable in abstract occupations than any other.

With these two results, there is more we can say about unobserved ability. First, as already mentioned, **Result 1** gives upper bounds for each unobserved ability. Second, **Result 2** provides another set of bounds – either upper or lower bounds, depending on each task. Third, **Assumption 2** requires the ratio of relative abilities to change always in a certain direction. This has very practical consequences. For example, consider workers in abstract occupations. **Assumption 2** implies that the "speed" at which these workers' unobserved routine ability changes as we move across i must not be

[44] A top lawyer might be a great taxi driver, maybe even better than many existing taxi drivers. Yet, the fact that most of these taxi drivers would be terrible lawyers makes the observed sorting optimal, without necessarily implying that the lawyer is worse at driving taxis than current taxi drivers.

[45] Notice that our model does not allow switchers to be in manual and abstract but not on routine tasks.

higher than the speed at which observed ability changes. This provides another set of bounds for unobserved ability. By combining all these sets of bounds, missing ability can be further narrowed down.

Even in this perfect scenario it is impossible to eliminate uncertainties in the imputation. This is partly due to the fact that AA refers to order among workers but not to levels or dispersion. Yet, the sole definition of ability as a continuous quantity denies any possibility of finding a unique imputation solution. In consequence, uncertainty in imputations cannot be eliminated. Naturally, the less gaps that are in Table 3.8, the more restrictions are generated on each blank cell, reducing the uncertainty. Yet, even in the limit of just one empty cell missing, uncertainty remains.

The uncertainty on finding estimates is even greater in a non-stylised setting where many factors affect wages, the market is far from perfect, workers do not have perfect information, and, importantly, absolute advantage does not hold. For example, later in Section 3.6, Figure 3.7 indicates that workers more or less cover the whole range of wages in all tasks, in contrast with the idealised labour market allocation of Figure 3.17. In consequence, impose restrictions on unobserved ability based on **Result 2** becomes unrealistic. A more practical approach is then needed, building upon the above results, but also allowing for some degree of realism.

In summary, even in a stylised setting, predicting population ability is an imprecise task. A rich set of boundaries can help lower this uncertainty, albeit never eliminating it.

Identifying absolute advantage

Given the relevance of AA in providing bounds for imputing unobserved ability, it is useful to develop a method to test this assumption. This test is fairly simple, and it is based on another result arising from **Assumption 1** and **Assumption 2**. This is:

Result 3. *Given a change in wage rates, those who move "upward" (i.e. from manual to routine, or from routine to abstract) have higher (lower) ability in their departure (destination) task, compared with those remaining (already) in the departure (destination) task. Similarly, those who move "downward" (i.e. from abstract to routine, or from* *routine to manual) have lower (higher) ability in their departure (destination) task, compared with those remaining (already) in the departure (destination) task.*[46]

To develop a useful test, it is revealing to contrast the case of AA with the alternative of non-hierarchical sorting. Under the latter, there is complete specialisation among workers and tasks. As such, the best workers are always allocated to each task. In terms of Figure 3.17, the coloured areas are always to the right of the distribution (as for abstract). This contrasts with AA, where the perfect match occurs only in one task – abstract, and the sorting is hierarchical.

Of course, **Result 1** does not change under non-hierarchical sorting. **Result 2** does. More precisely, those who are never observed in a given occupation have lower ability than all those who are ever in that occupation. Yet, since this result relies on unobserved ability, it cannot provide a test for AA. Luckily, **Result 3** is also different. Since switchers always come from and enter into the bottom of those employed in an occupation, the ability of switchers compared with those in the origin and destination occupations is **always** lower. This is a sharp contrast with AA, where ability of switchers might be lower or higher than the alternative comparison group. Since **Result 3** relies only on observed ability, it provides a useful test to evaluate the extent of AA in the dataset.

3.8.2 Model evaluation

To justify the usage of the above model to estimate unobserved ability, we can proceed to test the key model assumptions.

Do workers optimise?

Recall **Assumption 0**: (i) the labour market is competitive; (ii) workers make optimal occupational decisions based on the current period's wage; (iii) there is perfect information on wage rates; (iv) there is full mobility between tasks. Under this circumstances, occupational decisions should reflect an optimisation process by workers. This is the

[46]Since this result holds for all switchers, it also holds when studying average ability of switchers relative to average ability of non-switchers. This is a more practical approach when using data.

Result 1 highlighted in the previous section. The dataset allows us to evaluate this result by looking at switchers. Recall that for this group, ability has been estimated for at least two occupations. Since wage rates are also known, a simple test for optimisation can be constructed by comparing the wage a worker obtains in the "current" occupation with that which he or she would obtain if were to *immediately* switch to the "alternative" occupation at which he or she are observed. If the difference is positive, it means the worker would be better off by switching.

From the log of equation (3.14), and given the set of assumptions just mentioned, it must be true that:

$$(\ln \omega_{kt} - \ln \omega_{jt}) + (\ln \eta_{ik} - \ln \eta_{ij}) < 0 \quad \forall t \tag{3.16}$$

where j refers to the occupation in which a worker currently is, and k refers to any other occupation in which the worker has been or will be in the future. If a worker is better of as his or her current occupation, the expression in equation (3.16) must be always negative.[47]

Results for this test are presented in Table 3.9. Each cell presents the average of equation (3.16) within switcher types. To facilitate analysis, these values were divided by each individual's log wage. As thus, they represent **percentage change relative to current wage**. For example, on average, all those who are currently in a routine occupation and have ever switched from or will switch to an abstract occupation, would see a increase in the log of their real wage equivalent to 1% if they were to immediately switch to an abstract occupation. Of course, this evaluation is valid only in a strong *ceteris paribus* sense. This is, as if just one worker switched at a time (otherwise, general equilibrium aspects would come into play). As the table reveals, not all workers might be optimising. On average, moving "upwards" would be profitable, whereas moving "downwards" would not be. Nevertheless, these values are relatively small. For example,

[47] Notice this is a *ceteris paribus* exercise that assumes the labour market is "dense" enough for a worker being able to switch occupations whilst keeping all other job characteristics constant. In reality, even if workers' switching is solely motivated by changes in wages, switching might not be *ceteris paribus*. For example, workers might also change industry or region. Wage differentials would need to be taken into account for a precise analysis. Introducing this into the analysis however is not possible, since there is no a-priori method to tell which jobs characteristics would workers trade when changing occupation.

Table 3.9: Average percentage change in log real wage from an immediate change in occupation, switchers only, HT estimation

		Alternative task		
		Abstract	**Routine**	**Manual**
Current task	**Abstract**		-1.4	-1.5
	Routine	1.0		-1.1
	Manual	1.4	0.9	

Table 3.10: Percentage of switchers for which immediate change in occupation would imply an increase in real wage, HT estimation

		Alternative task		
		Abstract	**Routine**	**Manual**
Current task	**Abstract**		34.0	37.0
	Routine	60.3		38.6
	Manual	59.9	54.5	

the average log real wage level for manual tasks over the period is around 6.5 (Figure 3.6). In contrast, on average, workers that would have moved from manual to routine would have increased their total wage by just 0.9%. This is arguably a negligible amount. This table might also reveal a dichotomy between those who move up as part of their career progression – expected to make gains, and those who move down due to reallocation, lay-off, or other involuntary factors. In other words, this table might be capturing some selection between switchers.

To explore these issues further, Table 3.10 presents the proportion of workers who would be better off by switching to an alternative occupation. Moving "upwards" would be profitable for more than half of switchers. Interestingly, there is still a third of switchers for which an immediate switch "downwards" would improve the real wage.

A less strict evaluation would use the lag of the wage rate in the alternative occupation (i.e $\ln \omega_{k,t-1}$ in equation 3.16), in order to allow for delays in workers' learning about pay on each occupation. Using this reduces each proportion in Table 3.10 by around 5 percentage points. However, recall wage rates are increasing over the period (Figure 3.10), making the above result quite mechanical. Another relaxation of the model would consider a higher threshold needed to motivate switching. This might be the case because of mobility costs, compensation for loss of tenure, and firm, industry or

Table 3.11: Percentage of switchers for which immediate change in occupation would imply at least a 5% increase in real wage, HT estimation

		Alternative task		
		Abstract	**Routine**	**Manual**
Current task	**Abstract**		3.0	3.4
	Routine	10.8		6.3
	Manual	16.3	15.7	

national wage increases. Table 3.11 presents an equivalent to Table 3.10, but using a 5% increase in wage as condition to move. The proportion of beneficiaries from switching is considerably lower than before.

The analysis so far mixes those who were previously in the "alternative" occupation but are now in the "current" occupation with those who already switched to the "alternative" occupation. A deeper analysis would look into the switching dynamics by differentiating between these two groups. Although not shown, the pattern is the same. Namely, on average, for those who are about to switch up (in the coming periods) it is profitable to do so, whereas for those who are about to switch down it is not. Likewise, for those who have already switched up it is not convenient to switch down again, whereas for those who have switched down it would be preferable to move up again. In other words, it is not a matter of having already switched or not, but about the direction of switching. The absolute numbers are also informative. There are around 16,000 instances (individual-year) involving "upward" switching, compared with 14,000 instances of "downward" switching. Many of the latter, strictly speaking not rational (in our simple model), are still a sizeable proportion of overall switching.

In summary, a strict version of **Result 1** needs to be rejected. More precisely, why do no more workers switch upward, if it is apparently convenient for them to do so? Similarly, why are many individuals switching downwards, when it seem irrational to do so? As already highlighted, for a great number of these switchers, the gain would be smaller than 5%. In a world with mobility costs, non-pecuniary factors, imperfect information, and so on, these numbers might still reflect an important degree of optimisation among the workforce.

Table 3.12: Pairwise correlation between predicted ability, for switchers, HT estimation

	Manual	Routine
Routine	0.47	
Abstract	0.60	0.62

Absolute advantage

An central assumption of the model is absolute advantage in ability among workers (**Assumption 1**). A natural test for this is to look at the correlation of estimated ability between occupations. A high degree of AA implies the ordering of workers in terms of ability is quite similar across tasks. Notice however that comparative advantage imposes restrictions only to the *order* of workers across ability dimension, but is silent about their level and dispersion. In consequence, correlation is not a precise test for absolute advantage (e.g. AA might not imply correlation equal to one). Moreover, this method only focuses on switchers, for which information on ability for at least two tasks is present. In any case, this can still be an informative method.

Tasks' pairwise correlations are shown in Table 3.12, for the HT estimation (the correlations are around 5% lower for the FE model). First of all, they are all positive, rejecting a widespread alternative view of comparative advantage. In fact, if the degree of absolute and comparative advantage were fairly similar across the population, correlation would be much closer to zero. These correlations are fairly high too, in particular with respect to abstract tasks. This is consistent with the evidence summarised later, where the switching from manual to routine seems to contradict AA, whereas all other switching patterns support it.

A more detailed method to study the nature of absolute and comparative advantage is to look at the relative ability of switchers with respect to current workers in both the destination and the origin occupation. Recall that AA implies **Result 3**. This is, those who move "upward" have higher ability in their departure task and lower ability in their destination task, all relative to those currently are in a task. Similarly, those who move "downward" have lower ability in departure task and higher in destination task, relative to current workers. We have already seen evidence in support of this result with respect

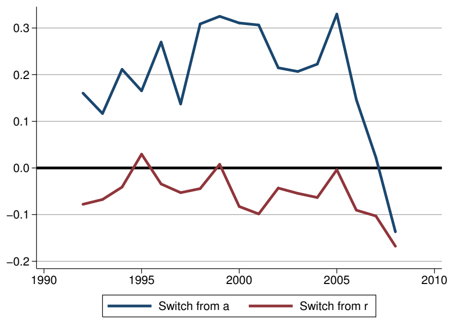

(a) Manual task

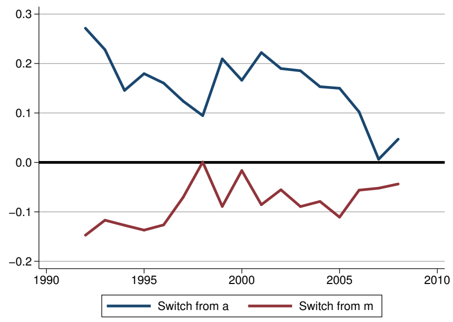

(b) Routine task

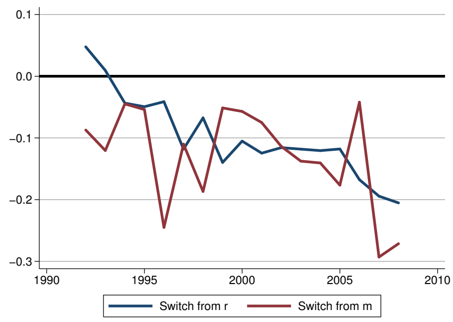

(c) Abstract task

Figure 3.18: Relative ability of “entrants” versus “incumbents” in a given task, HT estimation

to switchers’ destination (Figure 3.12). Figure 3.18 presents this evidence jointly for the three tasks, comparing “entrants” with “incumbents”. The clearest evidence of AA is on the routine task. For manual tasks, those coming from routine do not seem more able on average that incumbents, but those coming from abstract clearly are. In the case of the abstract task, those coming from both routine and manual seem clearly less able; but again, there is not much difference between the origin task.

There is also good evidence in favour of AA from the perspective of “leavers” versus “stayers”. Panel (a) shows the case of workers in manual occupations. Those who switch toward abstract tasks are in fact more able than stayers, strongly supporting AA. Yet, those switching to routine jobs are roughly equivalent if not worse than those who remain, playing against the model. This might explain the lower correlation between manual and routine ability in Table 3.12. Panel (b) greatly supports AA, as there is a clear distinction between moving up or down, out of routine tasks. Finally, Panel (c) shows that those switching away from abstract tasks are indeed less able than stayers. However, leavers are roughly equivalent in abstract ability, regardless of the destination.

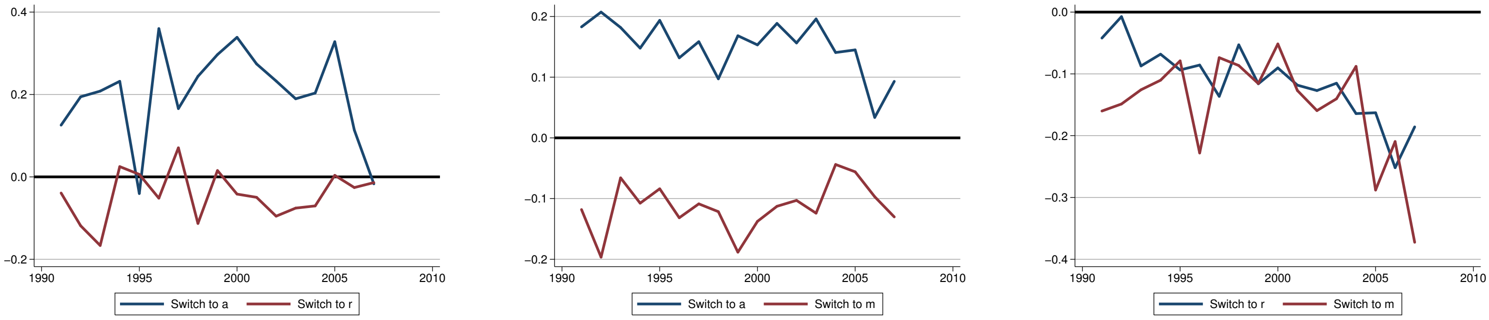

(a) Manual task (b) Routine task (c) Abstract task

Figure 3.19: Relative ability of "leavers" versus "stayers" in a given task, HT estimation

A story more consistent with AA would have shown differences between them.[48] All in all, this might not be such a serious issue, since according to Table 3.4, the transition from abstract to manual tasks is relatively low.

To conclude this section, there is an interesting degree of support to the rather simplistic assumptions of the model, in particularly after allowing for switching costs, which do exist in reality. Perhaps the largest departure from it refers to manual-routine absolute advantage. As said, this might be due to the inherent difficulty in classifying these tasks. As Bisello (2013) shows, there is an important degree of correlation between manual and routine components in many occupations.

3.8.3 Workers' ability in unobserved occupations

Section 3.7 characterised ability in occupations where workers *are* observed. Now it is time to impute ability for those occupations in which workers *are not* observed. Of the 15,843 individuals in the sample across the years, 68% are missing manual ability, 45% are missing routine ability, and 56% are missing abstract ability. Overall, 53% of ability data is missing, highlighting the relevance of this second step within the overall exercise. Namely, roughly half of the data is imputed using information from the other half.

In order to impute "unobserved" ability, I use two methods. **Method 1** is based on **Result 1**, which states that unobserved ability must be low enough in order for workers

[48]Notice however that our model in Section 2.3 does not allow switching between manual and abstract tasks.

not to switch to those occupations at which they are not observed. This imposes **upper bounds** on unobserved ability. There is already a natural bound at zero, so in effect this method involves, for each worker, a range of unobserved ability. The choice to make is then to select some point along these intervals.

Method 2 is based on **Result 2**, which states that unobserved ability must be strictly lower or higher than observed ability, but never in between (recall Figure 3.17). According to this result, there are distinctive ability intervals among tasks. It would then be possible to map each observed ability into any other task, leading to a more complete characterisation of the ability distribution. Now, the strictness of this result was already rejected. According to results in Figure 3.13, observed ability covers a considerable range, including values close to zero for all tasks, in contrast with the theoretical prediction just mentioned (again, see Figure 3.17). The solution used is to map observed distributions into each other in a less strict fashion, allowing for some overlapping.

Additionally, both methods are based on **Result 3**, which derives from absolute advantage. As such, both methods enhance the degree of absolute advantage of the final ability distribution with respect to those estimated from the data, and presented earlier in Table 3.12.

Method 1

First, we calculate bounds for missing ability, based on **Result 1**. Recall that for every worker and missing ability, the boundary chosen is the lowest among the period (see equation 3.15). These bounds define a complete interval, with zero at the lower end, in which "true" ability could reside. Now, we know that workers might not switch for small gains, because of switching costs or information problems. As such, these bounds might have a negative bias because the potential ability of workers in unobserved occupations could be higher than what a strict application of the model suggests. A rather trivial solution to the problem would be to simply assume the aforementioned boundary as missing ability. This is, assume unobserved ability is at the top of the interval. Certainly, this is likely to introduce a positive bias for most of imputed ability. Whether the two

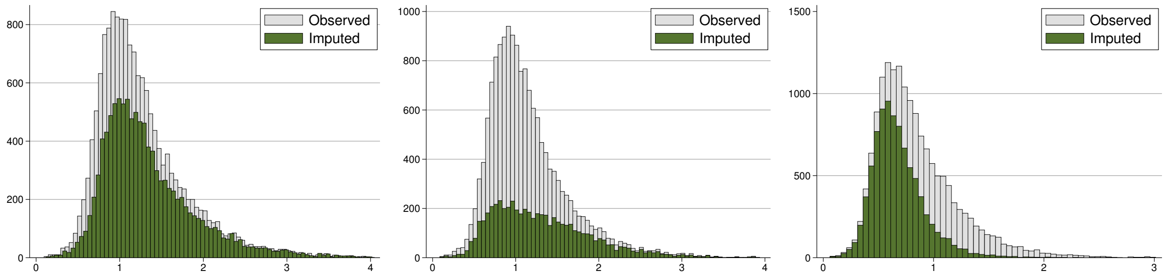

(a) Manual task (b) Routine task (c) Abstract task

Figure 3.20: Upper bounded ability distributions, HT estimation. The grey area (representing "observed" ability) is stacked above the green area (representing "unobserved" ability).

biases offset or not is impossible to tell. But this approach provides a first take on the possible population's ability distributions.

These distributions, called **upper bounded ability distributions**, are shown in Figure 3.20, for the HT estimation. These frequency graphs highlight both the observed (grey colour) and imputed (green colour) distributions of ability. The total distribution is computed by stacking these two distributions together. As the number of missing data suggested, manual distribution required the most of imputation (around two thirds), and routine the lowest. Interestingly, the positive skewness remains a feature of them all.[49] Importantly, this method also assumes absolute advantage. Since wage rates are equivalent across workers, it is only workers' observed ability that characterises each worker's boundary. Thus, ability differences between workers are mapped into other tasks. In effect, correlation between workers' ability across tasks is about 20 percentage points higher than the one presented in Table 3.7.

Instead of assuming the boundary, one could take unobserved ability to be the mean of the interval (i.e. half the boundary). This is also an arbitrary choice, but it is likely to reduce the positive bias introduced above. The result of this assumption is shown in Figure 3.21. There is still a consistent positive skewness among distributions. Correlation between abilities still goes up, but only by around 5%. The most salient difference between these two cases is the second spike for the abstract distribution. In

[49] A log-normal fit on each of them results in the standard deviation being 0.409, 0.403, and 0.421 for manual, routine, and abstract distribution respectively. The mean – recall – is of no valid interpretation. These numbers are useful for the simulations in Chapter 4.

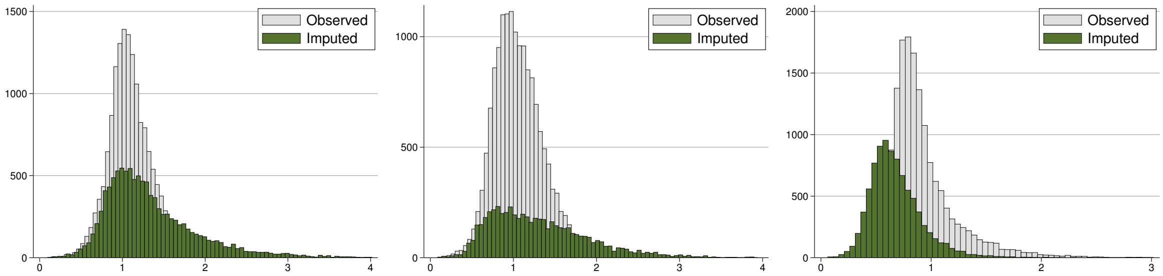

(a) Manual task (b) Routine task (c) Abstract task

Figure 3.21: Ability distributions assuming unobserved ability equals half of the upper boundary, HT estimation. The grey area (representing "observed" ability) is stacked above the green area (representing "unobserved" ability).

effect, the further away unobserved ability is assumed from the boundary, the more distinct these spikes are.

It might be evident by now that any other assumption regarding unobserved ability between half of the interval and the boundary will produce distributions that are in between the ones just presented. Considering the nature of the negative bias – due to switching and information costs, there is perhaps a greater case for unobserved ability to be between higher than predicted. Having this in mind, positive skewness of ability is likely to be a feature of the actual distributions.

Method 2

An alternative approach to find these distributions is based on a less strict form of **Result 2**. This form would not assume that **every** unobserved ability is lower or higher than an observed one, but rather than most of them are. As mentioned earlier, **Result 2** is largely based on AA, for which evidence was in general favourable. An example of this method goes as follows. To impute missing abstract ability, first find the mean of the observed abstract ability distribution (-0.03). Then, find the ability level of the worker at the 95$^{\text{th}}$ percentile of the observed routine ability distribution (0.47). Rescale the level of ability of routine workers such that the two numbers just found coincide. For all those missing abstract ability but not routine ability, the new rescaled routine level is used as the imputed abstract ability. There is a remnant of workers still missing abstract ability – those who have only been in the manual occupation. For these, repeat

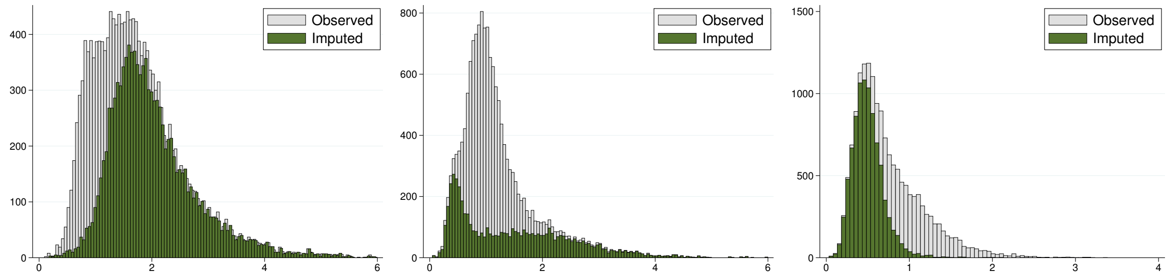

(a) Manual task (b) Routine task (c) Abstract task

Figure 3.22: Ability distributions obtained from merging observed distributions, HT estimation. The grey area is stacked above the green area.

the process, using the mean of the new combined abstract distribution as the merging point. A symmetric approach is followed to find the ability of all those missing manual ability. First, find the mean of the observed manual ability distribution (0.02). Then, find the routine ability of the 5th percentile observed routine distribution (-0.50). Now, rescale the latter. The new routine ability is the imputation for missing manual ability. The remaining can be mapped from abstract to manual as above. Regarding those missing routine ability, the process is the same, using the 5th and 95th percentile worker to map manual and abstract into routine, respectively.

The result of this imputation is presented in Figure 3.22. For the percentiles chosen, the final distributions are fairly consistent with the previous imputation method; all distributions are asymmetric with positive skewness.[50] Since this imputation maps the ordering of workers between tasks, it also enhances the resulting level of AA. Correlation between tasks increases by around 20 percentage points from those in Table 3.12. Just like before, the particular choice of merging points entirely defines the result. For instance, using the 90th (10th) percentiles instead of the 95th (5th) percentiles as merging points shifts the imputed and the complete abstract (manual) distribution to the right (left). On the contrary, using the 99th (1th) percentiles shifts the imputed and the total abstract (manual) distribution to the left (right). In the later case, the distributions have two spikes.

To conclude, ability distributions are likely to be positive skewed, in light of the relatively similar results produced by two different imputation techniques. Still, the difficulty of

[50] A log-normal fit finds the standard deviation to be 0.482, 0.542, and 0.527 for manual, routine, and abstract distribution respectively.

the task at hand does not allow for greater certainty about the counterfactual ability. Intuitively, the fact that a worker is not observed in a certain task might provide an idea of what ability level she or he **does not** have, but it might not reveal very much about what ability she or he **does** have. This uncertainty is intrinsic to the problem at hand, and to the best of my knowledge there is no model-based method where such uncertainty can be reduced, without imposing further restrictions on population's ability (as in Lemieux, 1998 based methods).

3.9 Conclusion

This chapter undertook the quest of identifying the ability distribution of the UK labour force, following an agnostic approach that imposes little structure on ability. Considering the intrinsic uncertainty of the problem, the results support a view of ability as having a positive skewed distribution, meaning an average ability level for the majority of the population on all tasks, with a minority being very skilled.[51] To the best of my knowledge, this is the first comprehensive attempt to characterise workers' complete ability distributions, filling an important gap in the literature.

Clearly, a way forward is to enrich the model to explicitly account for switching costs or information problems. Here, the treatment was more intuitive than anything. Also, workers occupational choices can be enriched by considering future period expectations of changes in wage rates. For instance, workers might consider changes in the present value of their income rather than current wage only. Adding occupation-specific experience and tenure might also enrich the sorting process.

A finer mapping between occupations and tasks might prove fruitful, in particular between manual and routine. Still, as evidence from Bisello (2013) indicates, a better refinement is not going to be easy. An m-to-1 mapping between tasks and occupations might also be beneficial. There is however the important barrier to identification that arises here. More time is expected to be devoted in the future to find a method that enables us to identify each ability from such a model.

[51] Recall the mean of the ability distributions is lost in the estimation. Hence, this analysis is true in *relative* terms, comparing only within tasks and not between them.

It might also prove useful to explore Lemieux, 1998 based methods to estimate ability. Here, all ability distributions were found to be positive skewed. Perhaps assuming equal distribution – the core of Lemieux's approach – might be reasonable, and certainly complementary to the results here.

Finally, the component which in my view would add the highest value added is introducing the firm to the dataset. Unfortunately, the BHPS data denies this option.[52] It is expected that in the near future this type of data will become more common. Alternatively, this study could be implemented for other countries like Germany, France, or the US, where firm-level data is already available. I expect my future research to focus in that direction.

[52]As said, Burdett, Carrillo-Tudela and Coles (2016) identify firms from workers' industry and occupational changes. For our purposes however, this method is not very helpful, because we have a third dimension – occupation. In effect, this method requires plenty of variation in firm/occupation to be able to truly identify worker-occupation and firm effects. This is, it requires workers to switch firm but not occupation, **and** occupation but not firm. Given the method they use to identify firms (change in occupation and industry), it is very likely that this variation is limited, hence bundling worker and firm heterogeneity together, as in our case. A truly useful mechanism would use firm-level data, where the sampling of several workers within firms makes the identification of worker and firm effect much more plausible. I expect to use such type of dataset in my future research.

4

Ability, technological change, and job polarisation in the UK

4.1 Introduction

Chapter 2 presented a simple model of tasks and sorting, which was then used to study job polarisation (JP). In particular, the model includes two independent sources or channels that might give rise to JP: (i) an increase in **workers' ability** at routine tasks, relative to their ability at manual and abstract tasks, and (ii) an increase in **productivity of routine inputs** in production of the final good, relative to abstract and manual inputs. These two assume $\sigma_Y < 1$, which this chapter shows is the most compelling empirical case. The objective of this chapter is to apply this model to the UK economy, in order to understand the role these two forces have played in observed polarisation in the UK.

The method to achieve this is relatively straightforward. From standard employment surveys (like the Labour Force Survey, or even the British Household Panel Survey), we observe the change in employment shares across occupations over the period of interest – 1992 to 2008. Since the two key endogenous variables in the model of Chapter 2 are the cut-off points defining workers' sorting into tasks (i_1 and i_2), these observed employment shares allow us to calculate the model equilibrium. Furthermore, Chapter 3 calculated the empirical distribution of ability for each task, for all the sample period. This is an important ingredient in the model. Importantly, the changes in ability distributions over time represent the supply side shocks that have affected the UK economy over the period. In turn, it is possible to characterise the labour demand shocks affecting the economy from these two components (employment shares, supply shocks) by means of a simple comparative static exercise. These demand shocks are those that reproduce the observed data.

Notice that this approach can only pin down biased task-biased technical change (TBTC), as opposed to any aggregate phenomena affecting the economy. Still, this is not a problem because what drives wage rates and sorting in the model is the change in relative productivities between tasks; Total Factor Productivity is irrelevant for sorting.

Once demand shocks are revealed, simple counterfactual exercises can help disentangling the role TBTC and skills have had on polarisation and inequality in the UK. Turning one of these off, and rerunning the model clearly shows that technology is the sole driver of both polarisation and inequality. Labour supply changes have played minor roles in sorting. If anything, they have led towards greater equalisation.

Finally, a preliminary exercise is conducted in order to shed some light on other countries' employment developments. For example, if the labour demand shock affecting the UK represents a relatively homogeneous process among advanced economies – for example due to a global technological process, it can be applied to other countries. Since the actual change in employment shares is observed, this can shed light on these economies' supply changes. Showing that these are highly unreasonable leads to the conclusion that the demand shocks affecting advanced economies, even if arising from a common source, have idiosyncratic consequences on each country. Naturally, to identify the nature of each country's labour demand changes we would need to conduct the analysis in Chapter 3 and this chapter for each country, clearly a burdensome and highly time consuming task.

This chapter contributes to the literature by providing the first attempt to evaluate the bias and magnitude of task-based technical change.

4.2 Fitting the model to the UK

The first step for estimating technological change in the UK is to "fit" the model to the UK economy. The use of the word "fit" instead of "calibrate" is deliberate. Here, the objective is not to find the equilibrium of the model – employment shares, but to find the "missing" (technology) parameters that make the model consistent with observed data. The only calibration conducted is that of the ability distributions, whose parameters can be deduced from the results in the previous chapter. In a sense, this is a "reverse" calibration exercise.

Recall from Chapter 2 that the solution of the model in any period involves the following two-equation system, which represent the demand and supply for relative tasks:

2x2 System

$$\left(\frac{T_r}{T_m}\right)^{1-\phi} = \left(\frac{\alpha_r}{\alpha_m}\right) exp\left[(\mu_r - \mu_m) + (\sigma_r - \sigma_m)\sqrt{2}\,\text{erf}^{-1}\left(2i_1^* - 1\right)\right] \tag{4.1}$$

$$\left(\frac{T_a}{T_r}\right)^{1-\phi} = \left(\frac{\alpha_a}{\alpha_r}\right) exp\left[(\mu_a - \mu_r) + (\sigma_a - \sigma_r)\sqrt{2}\,\text{erf}^{-1}\left(2i_2^* - 1\right)\right] \tag{4.2}$$

where T_j are given by

$$\begin{aligned}
T_m &= \Lambda_m\left[1 - \text{erf}\left(\frac{\sigma_m}{\sqrt{2}} - \text{erf}^{-1}\left(2i_1 - 1\right)\right)\right] \\
T_r &= \Lambda_r\left[\text{erf}\left(\frac{\sigma_r}{\sqrt{2}} - \text{erf}^{-1}\left(2i_1 - 1\right)\right) - \text{erf}\left(\frac{\sigma_r}{\sqrt{2}} - \text{erf}^{-1}\left(2i_2 - 1\right)\right)\right] \\
T_a &= \Lambda_a\left[1 + \text{erf}\left(\frac{\sigma_a}{\sqrt{2}} - \text{erf}^{-1}\left(2i_2 - 1\right)\right)\right]
\end{aligned} \tag{4.3}$$

and $\Lambda_j = \frac{1}{2}exp\left(\mu_j + \frac{\sigma_j^2}{2}\right)$. The solution of this system in terms of i_1^* and i_2^* represents the simultaneous equilibrium in the relative tasks markets, shown again in Figure 4.1.

This system of equations has three core elements. First, the equilibrium itself – i_1^* and i_2^*, related to employment shares through:

$$L_m^* = i_1^* \qquad L_r^* = i_2^* - i_1^* \qquad L_a^* = 1 - i_2^* \tag{4.4}$$

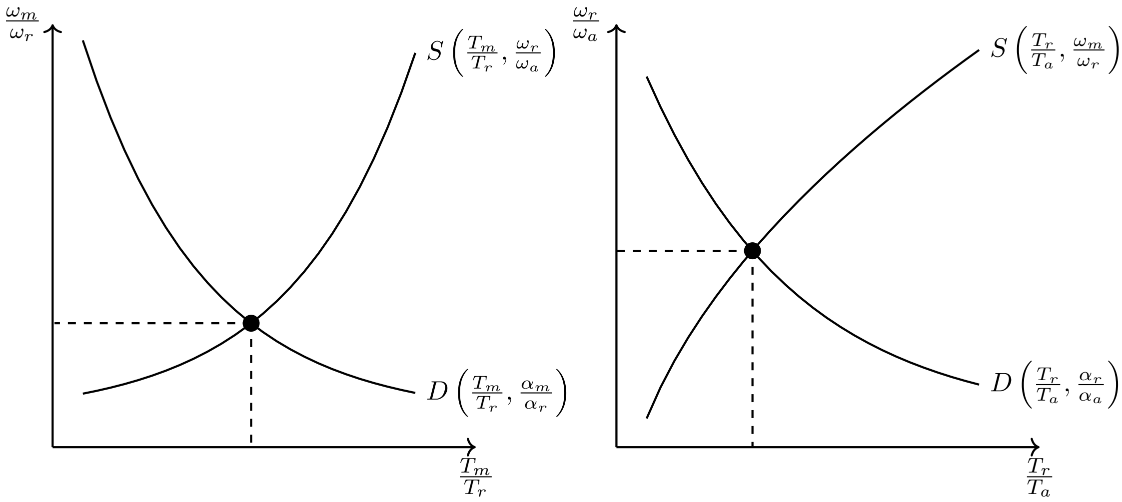

(a) Equilibrium in $\frac{T_m}{T_r}$ market (b) Equilibrium in $\frac{T_r}{T_a}$ market

Figure 4.1: Markets for relative tasks

Second, the supply of tasks, where the idiosyncratic parameters are those from the ability distribution, μ_j and σ_j. Third and final, the demand for tasks, where the idiosyncratic parameters are α_j and ϕ. This section evaluates the first two aforementioned elements. Next section focuses on the demand for tasks.

4.2.1 Employment shares

Employment shares per occupation – L_j in the model – are observed from any labour market dataset that includes occupation information. The only extra work required is to map occupational codes into tasks. This mapping has already been described in Section 3.6 of Chapter 3, and it is also used here. These shares fully determine the model's only endogenous variables, i_1^* and i_2^*, according to equation (4.4). Although the model and it's solution is written in terms of the latter two variables, here the focus is on L_j^*.

There are two data sources from where these shares are obtained. The first one is the British Household Panel Survey (BHPS), already used for the estimation of ability in Chapter 3. As such, it provides greater consistency with the ability parametrisation used here. Yet, the BHPS was shown to suffer from attrition. Thus, even if the initial sample (1991) was representative of the UK labour force, this qualification might have

Table 4.1: Labour market cut-off points, derived from observed employment shares for different years and samples

Sample	Emp. shares	Year 1992	1997	2008
BHPS	L_m	0.17	0.17	0.15
	L_r	0.51	0.52	0.44
	L_a	0.32	0.31	0.41
LFS	L_m	0.15	0.16	0.18
	L_r	0.51	0.48	0.41
	L_a	0.34	0.36	0.41

deteriorated over time. The second data source used is the quarterly Labour Force Survey (LFS). Although this is also a longitudinal survey, 20% of its sample renews every quarter, ensuring greater representation of the wider UK labour force over time.

Shares are calculated for three years: 1992, 1997, and 2008. 1992 is the beginning of the quarterly LFS survey, and also included in the BHPS; 1997 is the year since data for most of EU countries is available from Eurostat, as used in Chapter 3, hence allowing greater comparability across countries (useful for further analysis later on); 2008 is right before the Great Recession, avoiding potential cyclical confounding factors.[1]

Table 4.1 shows the outcome of this exercise, for selected years and samples. To complement this table, Figure 4.2 shows the absolute change in employment shares for each occupational group. Panel (a) is of course consistent with the evidence shown in Figure 3.9, in Chapter 3. There is no evidence of polarisation, but of a considerable shift from routine to abstract occupations, plus a slight fall in manual tasks. The majority of the change happens between 1997 and 2008 (judging by the comparison of patterns across periods). Conversely, the LFS dataset indicates the presence of polarisation, with a bias toward abstract tasks. Still, a great deal of this polarisation happens between 1997 and 2008, although some is observed between 1992 and 1997. The magnitude of the fall in routine tasks is not the same across surveys, but always lies between -0.07 and -0.09. Because the two data sources yield dissimilar employment patterns, both are used.

[1]Both the BHPS and the LFS define occupations based on the Standard of Occupations Classification (SOC). Regarding the specific period of the year used, the BHPS is an annual survey, yet the bulk of it is collected in the the last quarter of each year. For comparability, the fourth quarter of the LFS is used as well. Finally, the calculations from the LFS make use of the frequency weights included in the dataset, to improve its wider population representation.

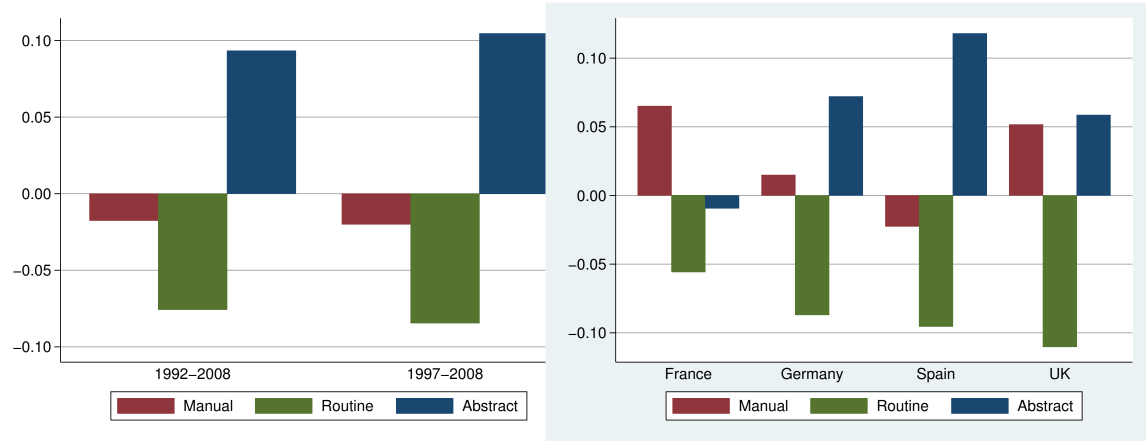

(a) BHPS dataset (b) LFS dataset

Figure 4.2: Absolute change in employment shares between selected years, in percentage points

4.2.2 Ability distributions

The parameters of the ability distributions can be obtained from the study presented in Chapter 3. There, the complete ability distributions were imputed for each task – manual, routine and abstract, using two methods. Since the model uses log-normal distributions, ability is taken from the estimations which more or less resemble these distributions.[2] In this chapter, **Method 1** refers to the estimates using the upper bounded ability distributions, and **Method 2** refers to the merged distributions using percentiles 95^{th} and 5^{th} as pivotal points. The resulting distributions for each of them are reproduced in Figures 4.3 and 4.4 respectively.

The next step is to parametrise each of these distributions, via maximum likelihood (e.g. Jenkins, 2004). More precisely, for a collection of data points $X = \{x_1, \ldots, x_N\}$ log-normally distributed with parameters μ and σ, these estimators are:

$$\hat{\mu}_{ML} = \frac{\sum_{i=1}^{N} \ln x_i}{N} \qquad \hat{\sigma^2}_{ML} = \frac{\sum_{i=1}^{N} (\ln x_i - \hat{\mu}_{ML})}{N} \tag{4.5}$$

This parametrisation is carried out for three particular years (see justification below). Changes in these distributions represent our supply shock. Results are presented in

[2]Naturally, the fact that the log-normal distribution is also used for the model's assumption is not a coincidence. They were chosen because the results in Chapter 3 highlighted positive skewed distributions. Among them, the log-normal is the most familiar example.

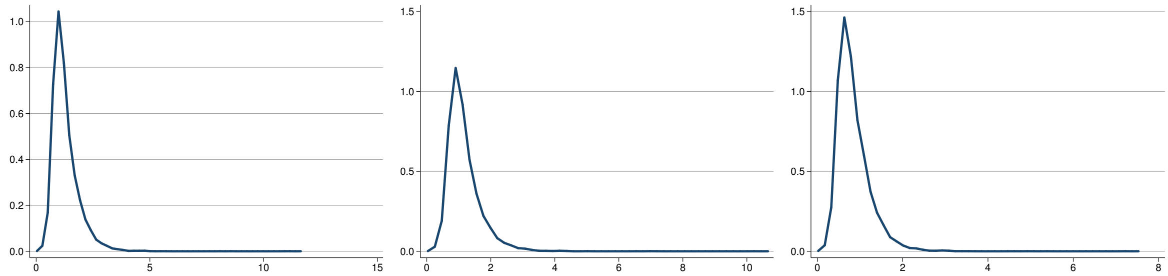

(a) Manual task (b) Routine task (c) Abstract task

Figure 4.3: Ability distributions, Method 1

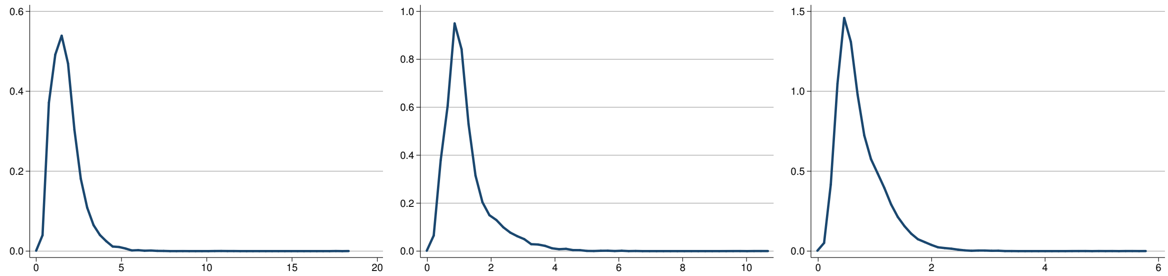

(a) Manual task (b) Routine task (c) Abstract task

Figure 4.4: Ability distributions, Method 2

Table 4.2. Both methods give statistically different values for each parameter and year. As such, using both methods might be informative. There is also significant variability in parameters over time. Allowing this variation to be reflected in the model's analysis is vital if we do not want to confound technological change with changes in ability when analysing job polarisation. The lack of interest in the empirical literature related to polarisation about the role of ability might introduce possible biases in the results. Interestingly, the σ_j parameters have shrunk around 10% over the 16 year period, indicating a lower degree of comparative advantage.[3] Additionally, there is positive support for one of the core assumption of the model, which is $\sigma_m < \sigma_r < \sigma_a$. In particular, $\sigma_m < \sigma_a$ in both method and for all years, whereas $\sigma_m < \sigma_r < \sigma_a$ holds for two out of six method-year estimations. This should not come as a surprise, given the strict nature of the model's assumptions. In particular, Chapter 3 shows that the support for absolute advantage between manual and routine tasks is lower than the evidence of absolute advantage with respect to abstract tasks. As such, there are workers

[3] Notice that σ_j^2 does not correspond to the variance of the log-normal distribution. This is $\left(exp\left(\sigma_j^2\right) - 1\right) exp\left(2\mu_j + \sigma_j^2\right)$. Still, a fall in one or both parameters leads to a fall in the variance of the distribution.

Table 4.2: Maximum likelihood estimates for log-normal ability distributions

Parameter	Year	Method	Task j Abstract	Routine	Manual
μ_j	1992	1	-0.147	0.115	0.243
		2	-0.277	0.173	0.55
	1997	1	-0.22	0.087	0.179
		2	-0.347	0.102	0.478
	2008	1	-0.261	0.076	0.137
		2	-0.364	0.08	0.425
σ_j	1992	1	0.423	0.401	0.411
		2	0.522	0.526	0.484
	1997	1	0.402	0.378	0.387
		2	0.50	0.492	0.464
	2008	1	0.376	0.346	0.357
		2	0.462	0.451	0.448

who the model would allocate into manual tasks but that are otherwise in routine jobs, and vice-versa. This is captured by the parametrisation shown in Table 4.2, since the violation of the core assumption is usually because $\sigma_r < \sigma_m$. Chapter 3 discusses several factors that might explain the latter result, including mobility costs between tasks, or imperfect information, or simply that absolute advantage is less prominent between these two tasks.

Finally, recall from Chapter 3 that we have not obtained an estimation of the **level** of ability, but only their relative values. As such, only σ_j and differences between μ_j are informative, but not the level of the latter. Yet, this does not pose a problem for our estimations because what determines the equilibrium allocation in the model is the relative level of abilities ($\mu_j - \mu_{-j}$), as equations (4.1) and (4.2) indicate. Recall that workers sort into occupation depending solely on their relative abilities between tasks, i.e. on comparative advantage (as equation 2.4 in the previous chapter indicates). Thus, an equal increase in ability across tasks does not alter allocation decisions.

4.3 Identifying technological change: methodology

4.3.1 Methodology

Having characterised the supply side of the model, and the actual equilibrium of endogenous variables, we can focus now on the demand side. Consider equations (4.1) and (4.2) again. Notice that each task output is a function of the ability of workers employed in that task, as per equation (4.3). As such, for known $\{i_1, i_2, \mu_j, \sigma_j\}$, every T_j is known. Similarly, the exponentiated terms in equations (4.1) and (4.2) are also known. What is left to identify are the technology parameters α_j and ϕ. Clearly, this is an undetermined two-equation system with four unknowns. Yet, as explained later, to capture technological change it is enough to identify the relative $alpha_j$s. For the moment, let us focus on the following two ratios: $\frac{\alpha_m}{\alpha_r}$ and $\frac{\alpha_a}{\alpha_r}$. Now, denote:

$$A = \frac{T_m}{T_r} \qquad B = exp\left[(\mu_r - \mu_m) + (\sigma_r - \sigma_m)\sqrt{2}\,\text{erf}^{-1}\left(2i_1^* - 1\right)\right]$$
$$C = \frac{T_a}{T_r} \qquad D = exp\left[(\mu_r - \mu_a) + (\sigma_r - \sigma_a)\sqrt{2}\,\text{erf}^{-1}\left(2i_2^* - 1\right)\right]$$

Then, the transformed system in equations (4.1) and (4.2) is:

$$\frac{\alpha_m}{\alpha_r} = f(\phi) \equiv A^{1-\phi}B \qquad (4.6)$$
$$\frac{\alpha_a}{\alpha_r} = g(\phi) \equiv C^{1-\phi}D \qquad (4.7)$$

This is a 2x2 system, where ϕ remains a "free" parameter, upon which the analysis can be conducted. In consequence, for a given parametrisation of the supply sector of the model and the employment equilibrium, we can derive a solution for the relative α_j ratios that depends **only** on ϕ.[4]

This methodology is sketched in Figure 4.5, using the example of the market for $\frac{T_m}{T_r}$. Panel (a) shows the "fitting" of the model presented in the previous section, where the

[4]As explained later, the freedom of σ_Y could be reduced by comparing predicted relative wages with the relative wages estimated in Chapter 3 (Figure 3.10). Unfortunately, due to apparent outliers, this is not a very helpful exercise. However, it reveals that the exercise is in fact quite accurate in matching empirical relative wages. See Section 4.4.3.

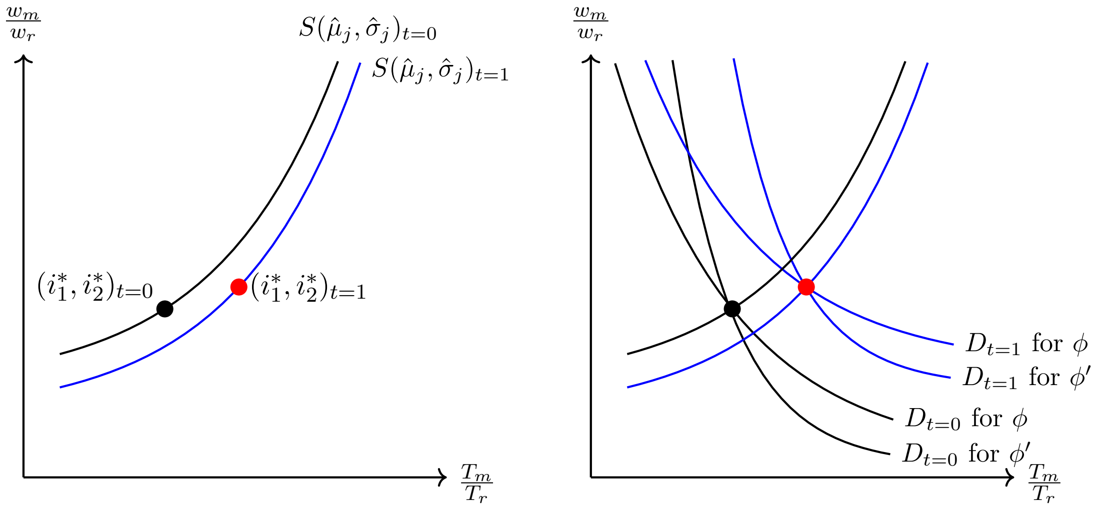

(a) "Fitted" equilibrium and supply of tasks (b) Identification of demand for alternative ϕ

Figure 4.5: Identification of demand in period $t = 0$ and $t = 1$, $\frac{T_m}{T_r}$ market

tasks' equilibrium and supply curves for two periods are identified.[5] From here, it is then possible to identify the demands, as Panel (b) shows. As said, these demands depend on the free parameter ϕ. Importantly, this comparative exercise is possible to carry out because there are no state variables like capital or workers' experience introducing dependence between periods. In consequence, any change in the model's equilibrium can be fully understood by evaluating it in the beginning and end of the period in question.

4.3.2 Modelling technical change[6]

The method outlined above can produce an estimation for $\frac{\alpha_m}{\alpha_r}$ and $\frac{\alpha_a}{\alpha_r}$, for any period, as a function of ϕ. The next step is to introduce technical change into the model, to which these ratios can be linked. The simplest approach to technological change is that of an aggregate process that increases total factor productivity (as in Chapter 1). This process is by definition unbiased, meaning all tasks are equally affected. In turn, the aforementioned ratios are not affected by such a process. This is, any neutral,

[5]This graph might give the impression that the supply is not required in order to identify the relative demand for tasks, if one has already characterised the model's equilibrium. This is not the case. The equilibrium in the tasks markets refer to $\frac{T_m}{T_r}$, which are a function of employment shares **and** ability parameters (see equation 4.3). Thus, both employment shares and ability are required to characterise the tasks markets' equilibrium.

[6]Part of the notation used in this section was already introduced in Chapter 2. Since the approach here is empirical, some differences exist. In particular, here the production function is normalised.

aggregate technical change affecting the UK economy in the period under study remains unidentified.

A more flexible specification assumes that exogenous technical progress affects the productivity of tasks' inputs T_j individually, – the so-called task-biased technical change (TBTC). This framework allows us to identify any **biased** technological progress, where biased refers to a change in the **relative** productivity of inputs. Any technical process that affects tasks equally does not alter relative productivities, and therefore is invisible to relative α_js.

Now, it is very important to notice that any analysis of ratios will not reveal the **level** of productivity of any input. Yet, they can reveal information about the **changes** in relative tasks' productivity. That is exactly the interest here, as TBTC refers to a process which alters tasks relative productivities. This is explained in detail later on.

In order to clarify exactly what we can measure, let us define two types of TBTC, which together completely characterise any possible pattern of TBTC. Given our particular focus on routine tasks – and the relevance they have for polarisation, the following categorisation is used:

- Routine-only task-biased technical change, RTBTC: it corresponds to a process that alters the relative productivity of routine tasks with respect to manual and abstract task, but **does not** alter the relative productivity between manual and abstract tasks.

- Non-routine-only task-biased technical change, NRTBTC: it corresponds to a process that alters the relative productivity between manual and abstract tasks.

This classification allow us to understand any TBTC as either RTBTC or NRTBTC; there is no third option. This is helpful because it enable us to think very precisely about the bias of the TBTC, which is in the end what enable its identification.

Now, let us retake the discussion about the ratios $\frac{\alpha_m}{\alpha_r}$ and $\frac{\alpha_a}{\alpha_r}$, which are obtained from the "fitting" exercise of above, given the free parameter ϕ. To see what these ratios can and cannot reveal about biased technical progress, individuals' tasks productivity needs to be added into the model. Now, since our focus is to compare estimation between two periods (e.g. 1992 and 2008), there is no need for a formal definition of technological change; the rate of growth in technology between these two periods can be simply encapsulated by the ratio of this task's productivity levels in these two periods. Thus, without loss of generality, let us denote $E_{j,t}$ as the productivity level of input T_j in period t. The "new" inputs in the production function are $E_{j,t}T_{j,t}$.

Additionally, according to Klump, McAdam and Willman (2012), for a comparative static exercise to be correctly defined it is appropriate to use a "normalised" production function. This is a function whose values at any point in time are relative to that of an arbitrary initial period. In our case, this normalised production function, with added $E_{j,t}$, is:

$$Y_t = Y_0 \left[\pi_{m,0} \left(\frac{E_{m,t}}{E_{m,0}} \frac{T_{m,t}}{T_{m,0}} \right)^{\phi} + \pi_{r,0} \left(\frac{E_{r,t}}{E_{r,0}} \frac{T_{r,t}}{T_{r,0}} \right)^{\phi} + \pi_{a,0} \left(\frac{E_{a,t}}{E_{a,0}} \frac{T_{a,t}}{T_{a,0}} \right)^{\phi} \right]^{\frac{1}{\phi}} \qquad t = 0, 1 \tag{4.8}$$

where $\pi_{j,0}$ corresponds to the factor's share in output in $t = 0$. Since this production function has constant returns to scale, $\sum \pi_{j,0} = 1$. A test for consistent normalisation confirms that $Y = Y_0$ for $t = 0$. Notice that this normalised function has not altered our original model. They are equivalent, with the mapping between specifications being $\alpha_{j,t} = \pi_{j,0} \left(\frac{Y_0}{T_{j,0}} \frac{E_{j,t}}{E_{j,0}} \right)^{\phi}$. All we have done here is to explicitly account for technological progress while providing a normalisation with respect to a baseline period.

Under the model's new specification, it holds that:

$$\frac{\alpha_{j,t}}{\alpha_{-j,t}} = \left(\frac{\pi_{j,0}}{\pi_{-j,0}} \right) \left(\frac{T_{-j,0}}{T_{j,0}} \frac{E_{-j,0}}{E_{j,0}} \right)^{\phi} \left(\frac{E_{j,t}}{E_{-j,t}} \right)^{\phi} \tag{4.9}$$

Notice we have estimated the left-hand side ratios from equations (4.6) and (4.7) **for any period**. Yet, we do not know any of the right-hand side terms. We can however find the rate of change of this ratio between two arbitrary periods. Using either 1992 or 1997 as our baseline period $t = 0$, and 2008 as $t = 1$, this is:

$$\frac{\left(\frac{\alpha_{j,1}}{\alpha_{-j,1}} \right)}{\left(\frac{\alpha_{j,0}}{\alpha_{-j,0}} \right)} = \frac{\left(\frac{E_{j,1}}{E_{-j,1}} \right)^{\phi}}{\left(\frac{E_{j,0}}{E_{-j,0}} \right)^{\phi}} \tag{4.10}$$

Thus, **the change in the α_js ratios provides an estimate of the relative rate of change of productivity between tasks**. This information can be used to achieve our goals, as explained below. But first, notice that the direction of change in productivity depends entirely on the sign of ϕ. Namely, the same change in relative α_js can lead to a relative fall ($\phi < 0$) or rise ($\phi > 0$) in relative productivities. In the pivotal case of $\phi = 0$ (i.e. a Cobb-Douglas production function), these ratios remain unidentified. This is the case because with these type of functions, biased and unbiased technical change are undistinguishable from the data. Namely, they can both be transformed into the other.[7]

To gain further knowledge about these ratios, let us introduce a useful definition of productivity $E_{j,t}$. This is:

$$E_{j,t} = \begin{cases} J, & \text{if } t = 0 \\ J + j, & \text{if } t = 1 \end{cases} \tag{4.11}$$

where $J > 0$. For example, productivity of routine tasks in the initial and final period is R and $R + r$ respectively. In this specification, technological change is revealed through variations in tasks' productivity. In our example, a technological change process affects routine tasks *iff* $r \neq 0$.

If we could identify the **absolute** productivity level of all tasks in any period ($E_{j,t}$), our characterisation of any technological process would be complete. We would know how technical progress has affected each task by simply looking at the change in their productivities. This identification is however not possible, as equations (4.9) and (4.10) reveal. In any case, notice we are interested in technological **change**. It turns out this is possible to understand solely from analysing changes in relative productivities. More precisely, these ratios reveal both the **bias** and the **magnitude** of this biased technological change. These are explained now.

4.3.3 Bias

Consider "routine-only" task-biased technical change, RTBTC. Without loss of generality, and using the notation in equation (4.11), let us define RTBTC as to $r \neq 0$ and $m = a = 0$.

[7]For example, biased technical change as in $Y = (AL)^{\alpha} K^{1-\alpha}$ is equivalent to unbiased technical change as in $Y = B\left(L^{\alpha} K^{1-\alpha}\right)$, with $B = A^{\alpha}$.

Then, the **percentage** change in productivity ratios are:[8]

$$\frac{\frac{E_{m,1}}{E_{r,1}} - \frac{E_{m,0}}{E_{r,0}}}{\frac{E_{m,0}}{E_{r,0}}} = \frac{\frac{M}{R+r} - \frac{M}{R}}{\frac{M}{R}} \qquad \text{and} \qquad \frac{\frac{E_{a,1}}{E_{r,1}} - \frac{E_{a,0}}{E_{r,0}}}{\frac{E_{a,0}}{E_{r,0}}} = \frac{\frac{A}{R+r} - \frac{A}{R}}{\frac{A}{R}}$$

Both M and A cancel out, resulting in an identical percentage change in both ratios, equivalent to $\frac{-r}{R+r}$. Thus, this is the case of a symmetric bias with respect to manual and abstract occupations. For instance, a fall in the **absolute** productivity of routine tasks ($r < 0$) leads to an increase in the **relative** productivity of manual and abstract tasks of the same magnitude. If tasks are highly substitutable ($\phi > 0$), this expands the relative demand for manual and abstract tasks, increasing the relative wage rates of them against the routine tasks. This changes worker's incentives, leading to reallocation. More precisely, i_1 increases and i_2 reduces, thereby augmenting L_m, L_a, and contracting L_r; this is, job polarisation. As such, RTBTC is conceptually the simplest technological change leading to job polarisation, a result which was already presented in Chapter 2.[9]

In the case of "non-routine-only" technical change (NRTBTC), there is a change in at least one other task's absolute productivity apart from routine. For instance, consider the case of a variation from R to $R + r$ and from A to $A + a$, with $m = 0$. Then, the change in the respective ratios is:

$$\frac{\frac{E_{m,1}}{E_{r,1}} - \frac{E_{m,0}}{E_{r,0}}}{\frac{E_{m,0}}{E_{r,0}}} = \frac{\frac{M}{R+r} - \frac{M}{R}}{\frac{M}{R}} \qquad \text{and} \qquad \frac{\frac{E_{a,1}}{E_{r,1}} - \frac{E_{a,0}}{E_{r,0}}}{\frac{E_{a,0}}{E_{r,0}}} = \frac{\frac{A+a}{R+r} - \frac{A}{R}}{\frac{A}{R}}$$

which can be reduced to:

$$\frac{-r}{R+r} \qquad \text{and} \qquad \frac{a\frac{R}{A} - r}{R+r}$$

Clearly, they are different whenever $a \neq 0$. For example, technical progress could lower absolute productivity of routine tasks in producing output ($r < 0$) whilst increasing absolute productivity of abstract tasks ($a > 0$). The result is the same as above, in the sense that the relative productivity of manual and abstract tasks increases with respect to routine tasks. Yet, the increase for abstract is larger than for manual tasks. In this

[8]It was mentioned above that we can obtain an estimation of $\frac{E_{j,t}}{E_{-j,t}}$. This is in fact unnecessary. By calculating the ratio of the relative α_js from equation (4.9) (after taking their *phi*-root), it is possible to find the rate of change of relative ability, as required here. In practice, this is the approach used.

[9]If tasks' substitution is relatively low ($\phi < 0$), the aforementioned technological changes leads to an increase in routine employment at expense of manual and abstract occupations. Therefore, to produce job polarisation, what is needed is an improvement in routine tasks' productivity relative to the other two.

case, job polarisation can also arise. Yet, for a high enough, employment in manual occupations could actually fall.

What is clear from this exercise is that we can learn about the **bias** of the technological process by analysing the proportional changes in relative productivities between two periods. If they are roughly equivalent, technological change is only affecting routine tasks – in relative terms. If they are not, the bias is indicated by their relative values.

4.3.4 Magnitude

These ratios also provide information about the **magnitude** of this technological change. Consider the case of RTBTC, as defined before. For an initial productivity level of R, an equal change to r would lead to a change in relative productivities equal to $\dfrac{-r}{R+r}$. This can be re-written as $\dfrac{-1}{\frac{R}{r}+1}$. So, from an empirical evaluation of this ratio we can obtain a measure of $\frac{r}{R}$. This reveals the magnitude of the technological progress affecting routine tasks alone, without need to pin down R and r. Naturally, this magnitude is consistent with several productivity change patterns. For example, RTBTC can occur from: (i) an increase in the **absolute** productivity of routine tasks only, (ii) a fall in the **absolute** productivity of manual and abstract tasks only, (iii) a combination of the two. With our method it is impossible to identify which one is happening in the data. However, that is not a problem. The magnitude of the change in relative productivities is still revealed.

To see this more clearly, notice that to identify the size of the technological process is enough to find, for example, the relative values of R and r, as in $\frac{r}{R}$. This ratio indicates the **change** in the **absolute** productivity level of routine tasks, R. That is our working definition of technological change. Using the RTBTC example, we found that this leads to a change in the relative productivity of other tasks of magnitude $\dfrac{-r}{R+r}$, which can be rewritten as $\dfrac{-\frac{r}{R}}{1+\frac{r}{R}}$. So, from an empirical evaluation of this ratio we can obtain a measure of $\frac{r}{R}$. Similarly with in our previous example of NRTBTC, it is enough to find $\frac{r}{R}$ and $\frac{a}{A}$ to characterise the process affecting tasks. Since the effect of technical progress is asymmetric, there are two equations from where these two factors can be identified. Re-written in a useful way, they are:

$$\frac{-\dfrac{r}{R}}{1+\dfrac{r}{R}} \quad \text{and} \quad \frac{\dfrac{a}{A}-\dfrac{r}{R}}{1+\dfrac{r}{R}}$$

From these two we can identify both $\frac{r}{R}$, and $\frac{a}{A}$.[10] Notice that first we must evaluate the bias in order to decide how many productivity terms to find out.

In summary, to characterise the technological force affecting the UK economy between two points in time, it is enough to use equations (4.6), (4.7) and (4.10), together with data from tables 4.2 and 4.1 as inputs.

Before moving on to the results, notice that this comparative statics exercise aiming to identify TBTC already takes into account any variation in the distribution of abilities of the population in the two periods. This is achieved by means of using year-specific μ_j and σ_j, as in Table 4.2. In other words, **this exercise is not confounding the sources of variations in employment shares** (assuming our model is "correct"). Naturally, it is of interest to disentangle the exact effect of both demand (technology) and supply (ability) on polarisation, as done later.

4.4 Identifying technological change: results

The only remaining uncertainty when estimating technological change is ϕ, or alternatively, the elasticity of substitution between tasks, equal to $\sigma_Y = \frac{1}{1-\phi}$. To narrow this down we can look into the literature. Barnes, Price and Sebastia Barriel (2008) estimates this elasticity for the UK, using firm-level data, finding that, for a variety of specifications and methods, this elasticity is never above 0.4. This is a relatively low level of substitution, and well below that of a Cobb-Douglas production function ($\sigma_Y = 1$), common in the theoretical literature. León-Ledesma, McAdam and Willman (2010)

[10]Since the change in productivity is relative, there is never need to model $\frac{m}{M}$, $\frac{r}{R}$, and $\frac{a}{A}$ together. At least one is redundant. For example, to obtain a symmetric relative increase in productivity of manual and abstract tasks with respect to routine tasks, these two are equivalent: (i) $m = a = 0$ and $r < 0$, (ii) $m > 0$, $a > 0$, and $r = 0$. Clearly, (i) is preferred. This redundancy highlights the fact that we cannot pin down the **absolute** level of productivity.

provide a brief summary of the literature regarding this estimation for the US, clearly picturing $\sigma_Y < 1$. The most recent study they refer to, Klump, McAdam and Willman (2007), finds an elasticity between capital and labour of around 0.56. Another recent study by Young (2013), also for the US, focuses particularly on individual industries. He finds that $\sigma_Y < 1$ for most of them, whereas $\sigma_Y = 0.62$ characterises the aggregate of them.

All these studies focus on the elasticity of substitution between labour and capital. In turn, they are not directly transferable into our setting without capital. It is not clear whether substitution between tasks is higher or lower than between capital and labour. There seems to be no study evaluating this. Thus, for the following results, I take a pragmatic approach by evaluating alternative values for σ_Y, around the ranges mentioned above. Importantly, ϕ is assumed to be constant throughout all the period.[11]

4.4.1 Bias

Table 4.3 shows the **proportional changes** in $\frac{E_{m,t}}{E_{r,t}}$ and $\frac{E_{a,t}}{E_{r,t}}$ between selected periods, for different values of σ_Y. The grey cells indicate the bias of the process, by highlighting the highest proportional increase (or in this case, the lowest fall), for each period, dataset and method of estimation of ability. First, notice that for $\sigma_Y < 1$, ratios are negative, implying the productivity of abstract and manual tasks has fallen **relative** to the routine occupations. This is expected. The idea behind a technological process that leads to polarisation, is that the productivity of routine tasks is increasing, for example, due to the fall in the price of capital (an alternative input to its production, e.g. Autor and Dorn, 2013). Because of the relatively low substitution between tasks ($\sigma_Y < 1$), employment shifts away from these more productive tasks in favour of less productive ones. The relative nature of this process needs to be stressed. There could have been an increase in absolute productivity of all tasks, but the increment in routine tasks' productivity is greater. Conversely, we arrive at exactly the opposite conclusion if we assume $\sigma_Y > 1$. In this case, productivity of tasks went up in both manual and abstract, *relative* to routine. More precisely, given high substitution, employment follows those

[11] Appendix F shows that any change in ϕ cannot by itself induce job polarisation, albeit it could have consequences in employment reallocation.

tasks that see relatively higher increase in their productivity. In turn, the employment dynamics seen in Figure 4.2 mean that productivity in abstract tasks has increased the most, and this increase is higher when using the BHPS data, leading to the observed fall in manual employment.

The results are qualitatively equivalent across parametrisations, namely the LFS (BHPS) dataset indicates technological change is biased in favour of manual (abstract) tasks compared with abstract (manual) tasks. These results make perfect sense with the change in employment shares observed in Figure 4.2. On the one hand, the largest relative fall in manual tasks' productivity (LFS) means more employment will shift toward them, to ameliorate their worse performance, given the relatively low substitution between tasks. In contrast, the BHPS estimates the largest relative fall in productivity to happen to abstract tasks, hence increasing employment in this tasks. The **relative** fall is high enough in order to lower employment in manual tasks too. This is the case in Panel (a) of Figure 4.2.

These results are not a coincidence, but rather mechanical. Recall that these relative changes are calculated **using the observed change in employment shares**. As such, they should reflect them. Still, this already highlights that changes in ability are not important enough to reverse these differences between surveys. More on this later.

It's worth noticing that the differences in the bias are lower for Method 1 than for Method 2. In light of the greater variance and wider mean ability levels coming from the latter – presented in Table 4.2, this is not a surprise.

4.4.2 Magnitude

Having established that the bias of technological change is non-symmetric with respect to manual and abstract tasks, let us proceed to evaluate the magnitude of this bias. Since our reference point is routine tasks, it is intuitive to make $r = 0$ and focus on the relative change in manual and abstract tasks' productivity. Therefore, here we want to identify both $\frac{m}{M}$, and $\frac{a}{A}$. According to the proportional change formulas, these are:

$$\frac{\frac{E_{m,1}}{E_{r,1}} - \frac{E_{m,0}}{E_{r,0}}}{\frac{E_{m,0}}{E_{r,0}}} = \frac{\frac{M+m}{R} - \frac{M}{R}}{\frac{M}{R}} \qquad \text{and} \qquad \frac{\frac{E_{a,1}}{E_{r,1}} - \frac{E_{a,0}}{E_{r,0}}}{\frac{E_{a,0}}{E_{r,0}}} = \frac{\frac{A+a}{R} - \frac{A}{R}}{\frac{A}{R}}$$

Table 4.3: Proportional changes in $\frac{E_{m,t}}{E_{r,t}}$ and $\frac{E_{a,t}}{E_{r,t}}$ between selected periods, alternative σ_Y

Survey	**Method**	**1992-2008** $\frac{E_{m,t}}{E_{r,t}}$	$\frac{E_{a,t}}{E_{r,t}}$	**1997-2008** $\frac{E_{m,t}}{E_{r,t}}$	$\frac{E_{a,t}}{E_{r,t}}$
LFS	1	-0.49	-0.43	-0.39	-0.35
LFS	2	-0.56	-0.45	-0.43	-0.36
BHPS	1	-0.16	-0.46	-0.23	-0.52
BHPS	2	-0.28	-0.48	-0.28	-0.53

(a) $\sigma_Y = 0.4$

Survey	**Method**	**1992-2008** $\frac{E_{m,t}}{E_{r,t}}$	$\frac{E_{a,t}}{E_{r,t}}$	**1997-2008** $\frac{E_{m,t}}{E_{r,t}}$	$\frac{E_{a,t}}{E_{r,t}}$
LFS	1	-0.65	-0.59	-0.53	-0.48
LFS	2	-0.72	-0.61	-0.58	-0.49
BHPS	1	-0.26	-0.62	-0.34	-0.67
BHPS	2	-0.40	-0.63	-0.41	-0.68

(b) $\sigma_Y = 0.6$

Survey	**Method**	**1992-2008** $\frac{E_{m,t}}{E_{r,t}}$	$\frac{E_{a,t}}{E_{r,t}}$	**1997-2008** $\frac{E_{m,t}}{E_{r,t}}$	$\frac{E_{a,t}}{E_{r,t}}$
LFS	1	-0.88	-0.84	-0.78	-0.74
LFS	2	-0.93	-0.85	-0.83	-0.74
BHPS	1	-0.48	-0.86	-0.58	-0.90
BHPS	2	-0.67	-0.87	-0.67	-0.90

(c) $\sigma_Y = 0.8$

Survey	**Method**	**1992-2008** $\frac{E_{m,t}}{E_{r,t}}$	$\frac{E_{a,t}}{E_{r,t}}$	**1997-2008** $\frac{E_{m,t}}{E_{r,t}}$	$\frac{E_{a,t}}{E_{r,t}}$
LFS	1	2.97	1.37	1.48	0.85
LFS	2	3.28	1.12	1.63	0.70
BHPS	1	0.45	1.44	0.39	1.85
BHPS	2	0.52	1.21	0.46	1.64

(d) $\sigma_Y = 1.4$

These ratios reduce to:

$$\frac{m}{M} \quad \text{and} \quad \frac{a}{A}$$

In other words, the results shown in Table 4.3 directly represent the relative change in productivities. First, as already highlighted, they are all negative, indicating a relative worsening with respect to routine productivity. Evidence for this was already found in Chapter 3, when estimating the wage rate of each task using the BHPS dataset (see Panel (b) of Figure 3.10 in the aforementioned chapter). Second, the lower the degree of substitution between tasks, the lower this relative worsening is. The intuition is the following. For a given change in employment shares to be explained, the more substitutable tasks are, the lower the adjustment in relative productivities required to produce that variation in employment shares. For the case of $\sigma_Y = 0.40$, the fall in abstract to routine productivity is around 0.44%, for a period of sixteen years. This implies a great improvement in routine technology over the period. Although this might sound like huge number, recall that this includes productivity changes due to higher capital stock used for produce routine tasks (as in Autor and Dorn, 2013).

It is also interesting to highlight that a great deal of the technological shocks affecting tasks occur between 1997 and 2008, judging by the comparison between periods. For example, 27% out of the 35% relative fall in abstract to routine productivity mentioned earlier happens between 1997 and 2008.

In summary, these results reveal a technological process that is strong and highly biased in favour of routine tasks, a phenomena highly consistent with the employment shifts observed in an economy with relatively low elasticity of substitution (as the evidence indicates). This process is strong even at very low substitution rates. The most trustworthy data source (LFS) implies this process is also more relatively biased towards abstract tasks (all this in relative terms). This is a sensible result too. The technological process affecting the UK economy has shown to be not only good for "robots" (which replace routine workers), but also relatively better for high-skilled occupations.

4.4.3 Relative wages and the elasticity of substitution

The above exercises considers σ_Y a "free" parameter. However, the simulation also produces a prediction of relative wages, based on Equation (2.27) (in Chapter 2), reproduced below:[12]

$$\frac{\omega_m}{\omega_r} = \frac{\alpha_m}{\alpha_r}\left(\frac{T_m}{T_r}\right)^{\phi-1}$$
$$\frac{\omega_r}{\omega_a} = \frac{\alpha_r}{\alpha_a}\left(\frac{T_r}{T_a}\right)^{\phi-1}$$

Since the analysis in Chapter 3 produced estimates for relative wages for each year, we can compare predicted and estimated wages in order to calibrate σ_Y. For instance, we can find the σ_Y which minimises the root mean squared deviation (RMSD) between simulated and estimated wage rates across the sample period.[13]

This exercise is shown in Table 4.4, where ability is parametrised using Method 1 and employment levels are taken from LFS. The first part of the table shows estimates for relative wages coming from Chapter 3 (see Figure 3.10), for selected years. The second part of the table shows predicted relative wages by the model for alternative values of σ_Y. The key here is the last column, which shows the root mean squared deviation of comparing predicted and estimated relative wages. As it is clear, the smallest is achieved for σ_Y is ∞. This is clearly an unrealistic case. The problem seems to be due to a relatively large divergence in the abstract wage. Empirically, this grew very fast towards the end of the period (2008), an increase that the simulated one cannot mimic; the estimated $\frac{\omega_a}{omega_r}$ is 1.4919, whereas the predicted is around 1.38 or 1.39. Actually, if the abstract wage is excluded from the RMSD (surely an arbitrary decision), the optimal σ_Y is 0.62.

Therefore, nothing much changes after performing this analysis. Yet, there are several important insights arising from this exercise. First, the RMSD is roughly 4 times higher when using ability estimated with Method 2 than when using ability estimated with

[12]Recall that this calibration exercise does not yield **level** of productivity. Therefore, individual wage rates cannot be computed. Hence, this analysis must be carried out in terms of relative wages.

[13]Recall from Chapter 3 that FE does not produce estimation of the level of wages, whereas HT does. Therefore, to compute relative wages, HT estimates are used.

		Estimated			
	Wages	1992	1997	2008	
	ω_m/ω_r	0.8756	0.8600	0.9299	
	ω_a/ω_r	1.3220	1.3655	1.4919	
		Simulated			
σ_Y	Wages	1992	1997	2008	RMSD
0.6	ω_m/ω_r	0.8858	0.9169	0.9452	0.0537
	ω_a/ω_r	1.2832	1.3428	1.3837	
0.8	ω_m/ω_r	0.8866	0.9178	0.9465	0.0531
	$\omega_a\omega_r$	1.2843	1.3440	1.3855	
1.4	ω_m/ω_r	0.8876	0.9189	0.9481	0.0523
	ω_a/ω_r	1.2858	1.3456	1.3880	
4	ω_m/ω_r	0.8885	0.9198	0.9496	0.0517
	ω_a/ω_r	1.2871	1.347	1.3901	
∞	ω_m/ω_r	0.889	0.9203	0.9503	0.0514
	ω_a/ω_r	1.2878	1.3477	1.3912	

Table 4.4: Comparison between estimated and simulated relative wages for alternative values of σ_Y, using Method 1 ability estimates and LFS employment data

Method 1, for the same σ_Y. Second, RMSD is marginally smaller when matching employment changes from LFS than BHPS. Third, for the parametrisation with the lowest RMSD (i.e. using calibrating employment shares with LFS data and ability distributions with Method 1), the match is quite good. For example, for $\sigma_Y = 0.6$, the percentage deviation between predicted and estimated wage ratios oscillate between -6% (underestimation) and 8% (overestimation) – a rather small error given such a simple model is used. Fourth, even if there is some error in the predictions, both the direction of wage changes (the sign) and the ranking of wages (the relative magnitude) is correctly predicted. All these results provide further support to the exercise.

4.5 Counterfactual exercises

Having characterised the biased nature of the technological change affecting the UK economy, several counterfactual exercises are now possible. These are very simple: we can "turn off" any of the two processes leading to changes in employment shares (changes in ability or TBTC), and see how employment shares would have behaved. Thus, the approach here is reversed. Instead of taking a labour market allocation from the data, we analyse how employment would have behaved under alternative circumstances. This requires a standard calibration and simulation exercise, the objective being to find the solution to the model's two endogeonous variables – i_1 and i_2.

4.5.1 Changes in ability only

The first counterfactual exercises assumes there was no TBTC affecting the economy. It enables to study the effect of changes in ability in employment developments. Recall from Table 4.2 that there are considerable changes in both the mean and variance of the log-normal ability distributions. In particular, μ_j become closer and σ_j falls. A way to test for this role is to "shutdown" the change in technology found in the previous section. This is, to assume $\frac{E_{j,t}}{E_{j,0}} = 1$. The rest of the calibration is taken from the ability parametrisation in Table 4.2. Importantly, the simulation takes employment shares in 1992 as the reference point, using the LFS dataset. Therefore, we assume $i_1 = 0.15$ and $i_2 = 0.66$ in the first period of the simulation (see Table 4.1). ϕ remains a "free" parameter.[14]

Figure 4.6 shows the counterfactual simulations using ability parameters from both methods 1 and 2 (see Table 4.2), for two calibrations of σ_Y. For better comparability, these graphs use the same vertical axis that those in Figure 4.2, to which these should be compared. There seems to be a mild fall in routine tasks using both methods and σ_Y calibrations. But the results for other tasks are inconsistent with each other, and if anything rather insignificant. Overall, it seems that ability developments have little effect on employment reallocation, let alone on polarisation. As the next counterfactual exercise shows, the polarisation story behind the LFS data seems to be essentially a technical change-driven process.

[14] The only caveat of this simulation exercise is the need to assume $\sigma_m < \sigma_r < \sigma_a$. Since this does not always hold in the data (in particular for ability derived from method 1), it is imposed by selecting σ_r to be marginally in between, thus changing its value as little as possible.

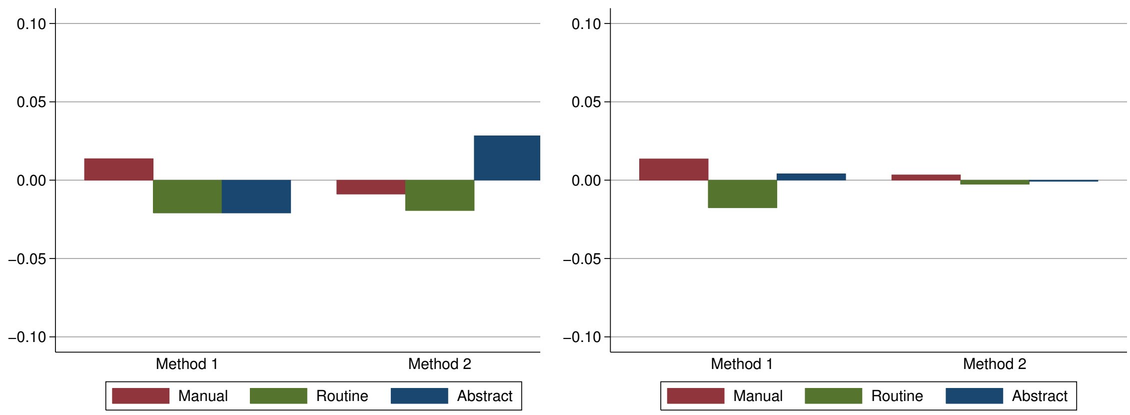

(a) $\sigma_Y = 0.4$

(b) $\sigma_Y = 0.6$

Figure 4.6: Absolute change in employment shares between selected years with respect to initial LFS shares, in percentage points, assuming no technological change

4.5.2 Technological change only

The alternative exercise is to consider the population as fixed between periods, and only see the consequences of technical change. This requires to input the results from Table 4.3 in the model, while using only the ability values in 1992. Again, the benchmark point corresponds to the employment shares from the LFS data. The results are presented in Figure 4.7, using both methods and two elasticities of substitution. Naturally, they complement the story told in the previous analysis. Polarisation arises quite distinctively, for both methods and reasonable σ_Ys. This is reflecting both the bias and strength of the technical change identified earlier. The magnitude of polarisation seem slightly larger than those in Panel (b) of Figure 4.2, perhaps reflecting some composite, ameliorating effect of technology and ability together.

4.5.3 Wage inequality and the middle class

While the analysis so far has focused on employment shares, the role of technology and ability on wage inequality is of particular interest. To measure inequality, we need to evaluate each worker's wage, which in turn requires an estimation of the wage rates ω_j. Since we cannot determine the level of productivity $E_{j,t}$, the **level** of wage rates is also not identified. Still, we can identify **relative** wages. This is enough if one is to calculate inequality using a scale-free index like Gini.

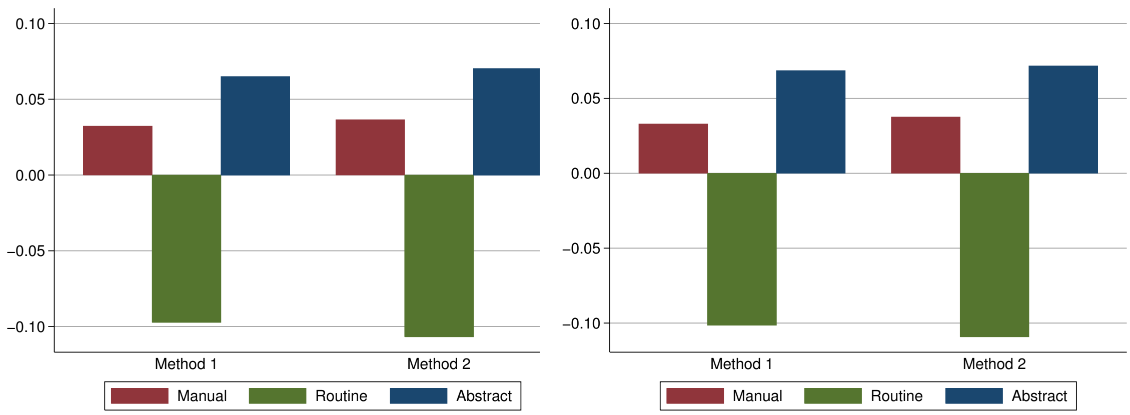

(a) $\sigma_Y = 0.4$ (b) $\sigma_Y = 0.6$

Figure 4.7: Absolute change in employment shares between selected years with respect to initial LFS shares, in percentage points, assuming fixed population

Combining equations (2.22) and (4.9) yields the following:

$$\frac{\omega_{j,t}}{\omega_{-j,t}} = \left(\frac{\alpha_{j,t}}{\alpha_{-j,t}}\right)\left(\frac{T_{j,t}}{T_{-j,t}}\right)^{1-\phi}$$

Since equations (4.6) and (4.7) allow us to estimate relative α_js for any period, and the tasks outputs are easily obtained from the labour market allocation, relative wage rates are identified. Thus, using our parametrisation of the ability distribution, and knowing the labour market allocation, either from Table 4.1 or from the counterfactual exercises above, we can evaluate the evolution of inequality under alternative scenarios. This is done by "populating" our model with an artificial dataset of N workers, so that they replicate the model's parametrisation and behaviour. Workers wages are calculated as the product of their ability and the wage rate at their chosen occupations (as in equation 2.26). From this set of wages, inequality is calculated.

Table 4.5 shows the results from several scenarios. First, it is interesting to bring in an external reference. Using quarterly LFS data, the Gini barely changed between 1992 and 2008 in the UK, remaining at 39 points. This is calculated using gross weekly pay. This contrasts with inequality in the model. The levels differ slightly, especially using the ability parametrisation from method 1. The model also indicates a considerable increase in inequality – between 13 and 20 points. Of course, a perfect fit is not expected. In our stylised model, no factor other than ability affects wages (not even hours worked). As the study in Chapter 3 highlighted, wage determination is a multidimensional phenomena.

Table 4.5: Evolution of Gini from LFS data, model, and counterfactual scenarios, alternative σ_Y

σ_Y	**Period**	**LFS**	**Model**		**Constant ability**		**Constant technology**	
			Method 1	Method 2	Method 1	Method 2	Method 1	Method 2
0.4	1992	38.6	44.6	37.1	44.6	37.1	44.6	37.1
	2008	38.9	57.7	57.6	60.8	60.3	37.9	31.2
0.6	1992	38.6	40.5	36.3	40.5	36.3	40.5	36.3
	2008	38.9	52.3	52.5	56.3	56.0	34.4	31.3

Note: LFS Gini is calculated using gross weekly pay ("grsswk" variable), with the dataset expanded using provided frequency weights. LFS income data is not available for 1992. Fourth quarter of 1993 is used instead.

Counterfactual exercises are quite revealing. Using our model as a benchmark, inequality would have decreased by 4 to 5 points if no biased technological change would have hit the UK economy over the period. The revelation is two-fold. First, technology is the sole driver of the observed increase in inequality in the model. Second, changes in ability lead to more equalisation over the period. This is because the ability distributions are becoming less heterogeneous. As mentioned, Table 4.2 reveals that the variance of each distribution is shrinking over time, whereas their mean differences are also reducing. Since wages map abilities, it is then not a surprise that inequality decreases in this scenario. Conversely, with a fixed population (that of 1992), inequality would increase even more than it does in the model. The story is the same for alternative values of σ_Y. Yet, in this case, the role of technology is lower. This was already mentioned when analysing Table 4.3, namely that higher substitution between tasks requires a technical progress of lower magnitude in order to explain a fixed degree of employment reallocation.

The analysis of the evolution of the middle class is also very interesting. This is presented in Table 4.6, where middle class is defined as a variable proportion of the population, in reference to a band **around the median wage.**[15] The result of the exercise are relatively similar to the one on inequality. In particular, there is a fall in the middle class under both definitions – larger if the interval around the median is wider. In the counterfactual scenario of constant skill distributions, the fall in the middle class

[15]Notice that defining the middle class as a fixed proportion is not reasonable, because the model only tells us about **relative** wages, and not about absolute wages, which is required if one were to calculate the absolute improvement in real income for a group of the population.

Table 4.6: Evolution of the middle class (as % of population), model, and counterfactual scenarios

σ_Y	**Period**	**Model** ± 25%	**Model** ± 50%	**Constant ability** ± 25%	**Constant ability** ± 50%	**Constant technology** ± 25%	**Constant technology** ± 50%
0.4	1992	41.8	51.0	41.8	51.0	41.8	51.0
	2008	38.0	41.0	35.2	41.2	44.2	49
0.6	1992	41.8	51.0	41.8	51.0	41.8	51.0
	2008	38.0	41.0	34.9	40.9	44.5	55.9

Note: the middle class includes all those with wages within a 25% or 50% band **with respect to the median**. Only Method 1 is shown, as results are quite similar.

would have been quite sharp, larger than when ability changes too (i.e. the model column). Conversely, turning off TBTC indicates that, the middle class would have gone up (except when $\sigma_Y = 0.4$, and the interval is wide). Clearly, their combination implies that ability has partially counteracted the strong effect of technological change in undermining the middle class in the UK, over the period.

Finally, the same result arises from an analysis of the (relative) income share of the top 1% and top 10% of the population. For example, in the model that assumes $\sigma_Y = 0.4$, the income share of these groups in 1992 is 3.8% and 26.8%. When both technology and changes in skills are considered, these shares increases to 4.1% and 30% respectively, in 2008. Without changes in skills, these shares would have gone up to 4.6% and 32.2% respectively. Conversely, without TBTC these would have gone down to 3.2% and 23%. The result is qualitatively equivalent for the case of $\sigma_Y = 0.6$. This is in line with the aforementioned analysis of the middle class. It confirms the pervasive effect that TBTC has had on polarisation, inequality, and the deterioration of the middle class in the UK.

4.5.4 Cross-country comparative analysis

As Figure 0.4 in the introductory chapter showed, polarisation patterns differ considerably across countries. A thorough method of evaluating cross-country differences would estimate wage determinants for each country, from where an ability distribution parametrisation can be obtained (essentially, replicate the analysis in Chapter 3 for other countries). Combined with data on employment shares, the bias and magnitude of the technological process driving employment reallocation can be assessed and compared. Clearly, this is an overwhelming task, and beyond the scope of this thesis – albeit an interesting window for future research.

An indirect approach would test the null hypothesis that the **bias** and **magnitude** of the technological shock affecting the UK economy over the period reflects a global force that plays out similarly in other countries. In the end, if technological change is driven by the emergence of China as a global manufacturing hub, the ICT revolution, the fall in trade costs, etc, it might be reasonable to expect shock to be similar across countries (i.e. a similar demand shock). If this is true, employment polarisation plus changes in ability distribution would map observed labour reallocation. Again, the problem is that ability is not observed. The only way around this is to start with certain values (e.g. those for the UK), and calculate how ability should have changed in order to reproduce observed polarisation. Recall that we do not need to worry about absolute ability parameters but their differences (e.g. $\mu_j - \mu_{-j}$). Thus, the dimensionality of the imputation is reduced by half. Still, the system remains underidentified, and extra assumptions regarding ability are required.[16]

It should be evident that countries following a relatively similar polarisation pattern than the one in the UK (e.g. Denmark, or Germany) would fare fairly well in this exercise. In the end, this is just a mechanical fit. Thus, an exercise on these countries is not very revealing. On the contrary, countries like France or Portugal show contrasting experiences. According to LFS Eurostat data, in these two countries reallocation of employment goes only toward manual occupations. In France, employment in routine and abstract occupations fell by 5.5 and 1 percentage points respectively between 1992 and 2008. Correspondingly, manual employment share increased 6.5 percentage points. In Portugal, the fall in routine and abstract occupations employment is around

[16]The precise method is the following. Assume the country under study follows the same distribution of abilities than the UK for 1992. Given an observed employment share, we can obtain $\frac{\alpha_{m,0}}{\alpha_{r,0}}$ and $\frac{\alpha_{m,0}}{\alpha_{r,0}}$ using equations (4.6) and (4.7) respectively. Second, assume the country faces the same technological force than the UK, from Table 4.3. Since the percentage change in α_js is directly related to this technological change (via equation 4.9), we can predict $\frac{\alpha_{m,1}}{\alpha_{r,1}}$ and $\frac{\alpha_{m,1}}{\alpha_{r,1}}$. Third, using equations (4.6) and (4.7) again together with observed i_1 and i_2 in 2008, we can find the ability parametrisation for 2008. Yet, there are two equations and four unknowns (relative μ_j and relative σ_j between two pair of tasks). The system is still undetermined. The extra assumptions imposed are that of keeping either of these as constant. This is clearly a simplification, but might still be informative of the plausibility of the null hypothesis.

Table 4.7: Required difference in ability distribution parameters to match observed employment shares, assuming common global technological force, for $\sigma_Y = 0.4$, method 1

Country	**Year**	$\mu_m - \mu_r$	$\mu_a - \mu_r$	$\sigma_m - \sigma_r$	$\sigma_a - \sigma_r$
UK	1992	0.128	-0.262	0.01	0.022
	2008	0.061	-0.337	0.01	0.03
France	1992	Assume same as UK		Assume same as UK	
	2008			-0.63	-0.64
	1992	Assume same as UK		Assume same as UK	
	2008	-1.63	-1.89		
Portugal	1992	Assume same as UK		Assume same as UK	
	2008			0.20	-0.09
	1992	Assume same as UK		Assume same as UK	
	2008	0.68	-0.47		

5 percentage points each.[17] Thus, these countries pose a major test to the hypothesis in question.

There is no need to dig deeper to reach some conclusions. Table 4.7 shows the required difference between parameters in order to match observed employment patterns. For reference, estimated values for the UK are included (taken from Table 4.2). The outputs of these exercises are shaded in grey. A quick comparison with those for the UK reveals they are largely out of scale. Certainly, it is unlikely to be true that initial ability distributions among countries are equivalent (as the "assume same as UK" cells presume). Yet, it is unlikely such differences are driven solely by these working assumptions. A more likely explanation is that the bias and magnitude of the change in relative tasks productivity affecting these two countries is different than the one affecting the UK. This might still be consistent with a global technological force. The difference is in how this force plays out among countries, and in particular how it affects relative demand for tasks. These differences might be due to capital quantity and quality, regulation, labour market flexibility, etc.

[17] It is worth noticing that the Eurostat dataset only categorises occupations up to one digit. This is a highly aggregated level, which contrasts with the four-digit level used for the UK. Accordingly, the mapping between occupations and tasks is imprecise, specially when distinguishing between manual and routine occupations. The former are exaggerated, as indicated by a comparison of results for the UK using locally sourced LFS versus Eurostat sourced LFS.

4.6 Conclusions

This relatively short chapter represents the culmination of the theoretical and empirical analysis presented in Chapters 2 and 3. By identifying the technological shock affecting the UK economy in recent decades, it demonstrated the pervasive role that trends in task-biased technical change had have in the UK labour market, leading to the well-known polarisation patterns observed in the dataset. Additionally, it was shown that TBTC tends to compress the middle class, while distribution greater share of total income to the top percentiles of the wage distribution.

The contribution of changes in ability to employment dynamics seems quite insignificant. Part of the reason seems to be because the differences between ability distribution parameters over time have not changed very much. If one were to see considerable polarisation, changes in the relative levels of ability might have been expected. In other words, over the period, any change in the skill of the UK workforce has been more or less homogeneous across tasks. The variations in absolute advantage needed to induce polarisation – or in fact any major reallocation – seems to be rather small, and clearly overshadowed by the magnitude of technical change. Still, it seems changed in ability acted as a check to changes in inequality and the compression of the middle class, albeit not high enough to stop their deterioration.

There is an important element missing in this analysis, mentioned earlier. This is the impossibility to distinguish between physical capital accumulation and TBTC *per se*. In our setting, capital is embedded in the relative α_j. This is enough for our purposes, where the focus is on employment and workers' wages, our final interest. But a more refined identification of TBTC would be ideal. Recall that the theory tells us that TBTC induces further capital accumulation (as in Autor and Dorn, 2013). As such, we can think of TBTC having a direct and an indirect effect of relative productivities. The approach here cannot disentangle between them. Still, the identification of this bundle is seen as a relevant addition to the literature, and certainly a starting point for further analysis.

5

Conclusion

This thesis has explored the role of technological change (both skill-biased and task-biased), investment in education, and skills in the rise and fall of the middle class, from both a theoretical and empirical perspective. There are many insights arising from the analysis, some of them reaffirming existing knowledge, and others bringing new contributions to the literature.

Among the most relevant contributions is the incorporation of a new channel that might give rise to polarisation – changes in workers' skills, filling an important gap in the literature. The simple model used to evaluate this phenomenon is also a contribution, and a strong starting point for further work. Some of the possible extensions have been highlighted already in the respective chapter. Although the results in the final chapter show that ability in the United Kingdom has not played a major role in polarisation, the story could be different in other countries; in particular in light of the dissimilar employment changes observed in other advanced economies (i.e. France) in the recent decades. This is certainly a hypothesis worth exploring. The methods this thesis provides could be a point of departure toward such goals.

The evaluation of the nature of task-based technical change (TBTC) is another contribution worth pointing out. The study of TBTC has been predominantly theoretical, besides its allegedly important effect on labour markets. This thesis provides direct evidence of that importance, at least for the United Kingdom. The results in this the

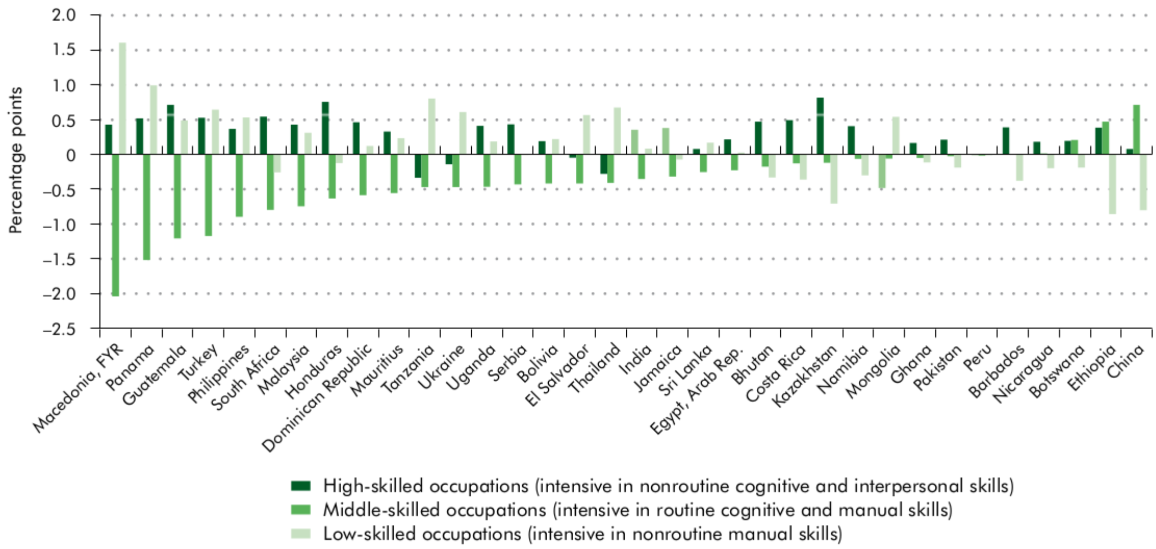

Sources: WDR 2016 team, based on ILO KILM (ILO, various years); the International Income Distribution database (I2D2; World Bank, various years); National Bureau of Statistics of China (various years). Data at http://bit.do/WDR2016-FigO_17.

Figure 5.1: Annual average change in employment share, between 1995 and 2012, from World Bank (2016)

last chapter also indicate that either advanced economies have not been affected by a common technological shock, or that this shock has affected them differently.

It might be also interesting to apply some of these methods to study the effect of technological progress, and TBTC in particular, in the labour market of developing economies.[1] For example, World Bank (2016) show that the labour market in many advanced economies has also polarised. Their main results in this respect is shown in Figure 5.1, which reproduces Figure 0.17 in the original paper. This raises interesting questions about the global nature of TBTC, and the role of education, and ability in general. In the figure it is also remarkable the case of China, which is that of "job compression". This might reflect strong interdependencies among countries, because of trade and global production networks. A theoretical and empirical attempt to connect job polarisation and job compression between economies is another interesting area of future research.

To conclude, I hope this thesis contributes toward a greater understanding of some issues of great importance today, and sheds more light on the circumstances that give

[1]There are greater challenges though, as the assumptions of perfect competition in labour markets and full mobility across occupations are less pertinent in these countries (e.g. Boeri, Helppie and Macis, 2008). The introduction of some rigidities might provide an insightful extension to our model.

rise to or can undermine the middle class, not only in advanced economies but, why not, also in developing ones.

Bibliography

Abowd, John M., Francis Kramarz and David N. Margolis (1999). 'High Wage Workers and High Wage Firms'. In: *Econometrica* 67.2, pp. 251–333.

Acemoglu, Daron and Joshua Angrist (2000). 'How Large Are Human-Capital Externalities? Evidence from Compulsory Schooling Laws'. In: *NBER Macroeconomics Annual* 15, pp. 9–59.

Acemoglu, Daron and David Autor (2011). 'Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings'. In: *Handbook of Labor Economics.* Ed. by David Card and Orley Ashenfelter. Vol. 4.B. Elsevier, pp. 1043–1171.

Acemoglu, Daron, Simon Johnson and James Robinson (2004). *Institutions as the Fundamental Cause of Long-Run Growth.* NBER Working Papers 10481. National Bureau of Economic Research, Inc.

Acemoglu, Daron and James A. Robinson (2002). 'The Political Economy of the Kuznets Curve'. In: *Review of Development Economics* 6.2, pp. 183–203.

Aghion, Philippe and Patrick Bolton (1997). 'A Theory of Trickle-Down Growth and Development'. In: *The Review of Economic Studies* 64.2, pp. 151–172.

Alesina, Alberto and Roberto Perotti (1996). 'Income distribution, political instability, and investment'. In: *European Economic Review* 40.6, pp. 1203–1228.

Alesina, Alberto and Dani Rodrik (1994). 'Distributive Politics and Economic Growth'. In: *The Quarterly Journal of Economics* 109.2, pp. 465–490.

Allison, Paul D (2002). *Missing data: Quantitative applications in the social sciences.* Sage Publications.

Alvaredo, Facundo, Anthony B. Atkinson, Thomas Piketty and Emmanuel Saez (2013). 'The Top 1 Percent in International and Historical Perspective'. In: *Journal of Economic Perspectives* 27.3, pp. 3–20.

Angrist, Joshua D. and Alan B. Keueger (1991). 'Does Compulsory School Attendance Affect Schooling and Earnings?' In: *The Quarterly Journal of Economics* 106.4, pp. 979–1014.

Arrow, Kenneth J. (1962). 'The Economic Implications of Learning by Doing'. In: *The Review of Economic Studies* 29.3, pp. 155–173.

Arteaga Cabrales, Carolina (2010). *Human Capital Externalities and Growth.* Borradores de Economia. Banco de la Republica de Colombia.

Asker, John, Allan Collard-Wexler and Jan De Loecker (2014). 'Dynamic Inputs and Resource (Mis)Allocation'. In: *Journal of Political Economy* 122.5, pp. 1013–1063.

Atkinson, A.B. (1983). *The Economics of Inequality*. Clarendon Press.

Atkinson, Anthony B. and Andrea Brandolini (2011). *On the identification of the "middle class"*. Working Papers 217. ECINEQ, Society for the Study of Economic Inequality.

Atkinson, Anthony B. and Salvatore Morelli (2014). *Chartbook of economic inequality*. Working Papers 324. ECINEQ, Society for the Study of Economic Inequality.

Autor, David and David Dorn (2013). 'The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market'. In: *American Economic Review* 103.5, pp. 1553–1597.

Autor, David and Michael J. Handel (2013). 'Putting Tasks to the Test: Human Capital, Job Tasks, and Wages'. In: *Journal of Labor Economics* 31.S1, S59–S96.

Autor, David, Lawrence F. Katz and Melissa S. Kearney (2006). 'The Polarization of the U.S. Labor Market'. In: *American Economic Review* 96.2, pp. 189–194.

Autor, David, Frank Levy and Richard J. Murnane (2003). 'The Skill Content of Recent Technological Change: An Empirical Exploration'. In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333.

Autor, David H. (2013). 'The "task approach" to labor markets: an overview'. In: *Journal for Labour Market Research* 46.3, pp. 185–199.

Autor, David H., David Dorn and Gordon H. Hanson (2013). 'The China Syndrome: Local Labor Market Effects of Import Competition in the United States'. In: *The American Economic Review* 103.6, pp. 2121–2168.

Baldwin, Richard (2013). 'Global supply chains: why they emerged, why they matter, and where they are going'. In: *Global value chains in a changing world*. Ed. by Deborah K. Elms and Patrick Low. Switzerland: World Trade Organization, Fung Global Institute, Nanyang Technological University.

Banerjee, Abhijit V and Esther Duflo (2003). 'Inequality and Growth: What Can the Data Say?' In: *Journal of Economic Growth* 8.3, pp. 267–99.

Banerjee, Abhijit V. and Esther Duflo (2008). 'What Is Middle Class about the Middle Classes around the World?' In: *Journal of Economic Perspectives* 22.2, pp. 3–28.

Banerjee, Abhijit V. and Andrew F. Newman (1993). 'Occupational Choice and the Process of Development'. In: *Journal of Political Economy* 101.2, pp. 274–98.

Barnes, Sebastian, Simon Price and Maria Sebastia Barriel (2008). *The elasticity of substitution: evidence from a UK firm-level data set*. Bank of England working papers 348. Bank of England.

Barro, Robert J. (2000). 'Inequality and Growth in a Panel of Countries'. In: *Journal of Economic Growth* 5.1, pp. 5–32.

– (2008). *Inequality and Growth Revisited*. Working Papers on Regional Economic Integration 11. Asian Development Bank.

Barth, Erling, Alex Bryson, James C. Davis and Richard Freeman (2016). 'It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States'. In: *Journal of Labor Economics* 34.S2, S67–S97.

Becker, Gary (1962). 'Investment in Human Capital: A Theoretical Analysis'. In: *Journal of Political Economy* 70.

– (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. University of Chicago Press.

Birdsall, N., C. Graham and S. Pettinato (2000). *Stuck in the Tunnel: Is Globalization Muddling the Middle Class?* Tech. rep. 14. Brookings Institution - Working Papers.

Bisello, Martina (2013). *Job polarization in Britain from a task-based perspective.Evidence from the UK Skills Surveys*. Discussion Papers. Dipartimento di Economia e Management (DEM), University of Pisa.

Blinder, Alan S. (2009). 'How Many US Jobs Might be Offshorable?' In: *World Economics* 10.2, pp. 41–78.

Boeri, Tito, Brooke Helppie and Mario Macis (2008). *Labor regulations in developing countries : a review of the evidence and directions for future research*. Social Protection and Labor Policy and Technical Notes 46306. The World Bank.

Bonjour, Dorothe, Lynn F. Cherkas, Jonathan E. Haskel, Denise D. Hawkes and Tim D. Spector (2003). 'Returns to Education: Evidence from U.K. Twins'. In: *The American Economic Review* 93.5, pp. 1799–1812.

Borissov, Kirill and Stéphane Lambrecht (2012). *The dynamics of income inequality in a growth model with human capital and occupational choice*. Working Papers. HAL.

Borjas, George J. (1987). 'Self-Selection and the Earnings of Immigrants'. In: *The American Economic Review* 77.4, pp. 531–553.

Bourguignon, Francois and Thierry Verdier (2000). 'Oligarchy, democracy, inequality and growth'. In: *Journal of Development Economics* 62.2, pp. 285–313.

Bourguignon, François and Christian Morrisson (2002). 'Inequality Among World Citizens: 1820-1992'. In: *American Economic Review* 92.4, pp. 727–744.

Bowles, Samuel and Herbert Gintis (2002). 'The Inheritance of Inequality'. In: *Journal of Economic Perspectives* 16.3, pp. 3–30.

Bryan, Lowell L. (2007). *The new metrics of corporate performance: Profit per employee*. URL: http://www.mckinsey.com/business-functions/strategy-and-corporate-finance/our-insights/the-new-metrics-of-corporate-performance-profit-per-employee (visited on 01/09/2016).

Brynin, Malcolm and Ayse Güveli (2012). 'Understanding the ethnic pay gap in Britain'. In: *Work, Employment & Society* 26.4, pp. 574–587.

Burdett, Ken, Carlos Carrillo-Tudela and Melvyn Coles (2016). 'Wage inequality: A structural decomposition'. In: *Review of Economic Dynamics* 19, pp. 20–37.

Card, David, Jörg Heining and Patrick Kline (2013). 'Workplace Heterogeneity and the Rise of West German Wage Inequality*'. In: *The Quarterly Journal of Economics*.

Card, David, Patrick Kline, Ana Rute Cardoso and Jorg Heining (2015). ‘Firms and Labor Market Inequality: A Review’. In: *Firms and the Distribution of Income: The Roles of Productivity and Luck*. National Bureau of Economic Research, Inc.

Carneiro, Pedro, Karsten T. Hansen and James J. Heckman (2003). *Estimating Distributions of Treatment Effects with an Application to the Returns to Schooling and Measurement of the Effects of Uncertainty on College*. NBER Working Papers 9546. National Bureau of Economic Research, Inc.

Caselli, Francesco and Jaume Ventura (2000). ‘A Representative Consumer Theory of Distribution’. In: *American Economic Review* 90.4, pp. 909–926.

Chamberlain, G. (1982). ‘Multivariate Regression Models for Panel Data’. In: *Journal of Econometrics* 18, pp. 5–46.

Choe, Chongwoo, Gloria Y. Tian and Xiangkang Yin (2014). ‘CEO power and the structure of CEO pay’. In: *International Review of Financial Analysis* 35.C, pp. 237–248.

Chudik, Alexander, M. Hashem Pesaran and Elisa Tosetti (2011). ‘Weak and strong cross-section dependence and estimation of large panels’. In: *The Econometrics Journal* 14.1, pp. C45–C90.

Ciccone, Antonio and Giovanni Peri (2006). ‘Identifying Human-Capital Externalities: Theory with Applications’. In: *The Review of Economic Studies* 73.2, p. 381.

Cortes, Guido (2016). ‘Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data’. In: *Journal of Labor Economics* 34.1, pp. 63–105.

Cortes, Guido M. and Giovanni Gallipoli (2014). *The Costs of Occupational Mobility: An Aggregate Analysis*. Working Papers 2014-015. Human Capital and Economic Opportunity Working Group.

Costinot, Arnaud and Jonathan Vogel (2010). ‘Matching and Inequality in the World Economy’. In: *Journal of Political Economy* 118.4, pp. 747–786.

Council of Economic Advisers (2015). *2015 Economic Report of the President*.

Dalmazzo, Alberto and Guido De Blasio (2005). *Where Do Human Capital Externalities End Up To?* ERSA conference papers ersa05p628. European Regional Science Association.

Deininger, Klaus and Lyn Squire (1998). ‘New ways of looking at old issues: inequality and growth’. In: *Journal of Development Economics* 57.2, pp. 259–287.

DeNavas-Walt, C., B. Proctor and J. Smith (2012). *Income, Poverty, and Health Insurance Coverage in the United States: 2011*. U.S. Census Bureau, Current Population Reports, pp. 60–243.

Dustmann, Christian and Maria Engracia Rochina-Barrachina (2007). ‘Selection correction in panel data models: An application to the estimation of females’ wage equations’. In: *Econometrics Journal* 10.2, pp. 263–293.

Dustmann, Christian and Nikolaos Theodoropoulos (2010). ‘Ethnic minority immigrants and their children in Britain’. In: *Oxford Economic Papers* 62.2, pp. 209–233.

Dutt, Pushan and Devashish Mitra (2008). 'Inequality and the Instability of Polity and Policy'. In: *The Economic Journal* 118.531, pp. 1285–1314.
Easterly, William (2001). 'The Middle Class Consensus and Economic Development'. In: *Journal of Economic Growth* 6.4, pp. 317–35.
– (2007). 'Inequality does cause underdevelopment: Insights from a new instrument'. In: *Journal of Development Economics* 84.2, pp. 755–776.
Ebenstein, Avraham, Ann Harrison, Margaret McMillan and Shannon Phillips (2014). 'Estimating the Impact of Trade and Offshoring on American Workers using the Current Population Surveys'. In: *The Review of Economics and Statistics* 96.4, pp. 581–595.
Ehrhart, Christophe (2009). *The effects of inequality on growth: a survey of the theoretical and empirical literature*. Working Papers. ECINEQ, Society for the Study of Economic Inequality.
Farzin, Y.H. (2009). 'The effect of non-pecuniary motivations on labor supply'. In: *The Quarterly Review of Economics and Finance* 49.4, pp. 1236–1259.
Felipe, Jesus and Franklin M. Fisher (2003). 'Aggregation in Production Functions: What Applied Economists should Know'. In: *Metroeconomica* 54.2-3, pp. 208–262.
Firpo, Sergio, Nicole Fortin and Thomas Lemieux (2011). 'Occupational Tasks and Changes in the Wage Structure'. IZA Discussion Paper No. 5542.
Forbes, Kristin J. (2000). 'A Reassessment of the Relationship between Inequality and Growth'. In: *American Economic Review* 90.4, pp. 869–887.
Fortin, Nicole and Thomas Lemieux (2016). 'Inequality and Changes in Task Prices: Within and between Occupation Effects'. In: *Inequality: Causes and Consequences*. Chap. 6, pp. 195–226.
Frydman, Carola and Dirk Jenter (2010). 'CEO Compensation'. In: *Annual Review of Financial Economics* 2.1, pp. 75–102.
Fuller, Matthew B. (2014). 'A History of Financial Aid to Students'. In: *Journal of Student Financial Aid* 44.1, pp. 42–68.
Gabaix, Xavier and Augustin Landier (2008). 'Why has CEO Pay Increased So Much?' In: *The Quarterly Journal of Economics* 123.1, pp. 49–100.
Gabaix, Xavier, Augustin Landier and Julien Sauvagnat (2014). 'CEO Pay and Firm Size: An Update After the Crisis'. In: *The Economic Journal* 124.574, F40–F59.
Gallup, J. (2012). 'Is there a Kuznets curve?'
Galor, Oded (2009). *Inequality and Economic Development: An Overview*. Working Papers 2009-3. Brown University, Department of Economics.
Galor, Oded and Omer Moav (2004). 'From Physical to Human Capital Accumulation: Inequality and the Process of Development'. In: *Review of Economic Studies* 71.4, pp. 1001–1026.
Galor, Oded, Omer Moav and Dietrich Vollrath (2009). 'Inequality in Landownership, the Emergence of Human-Capital Promoting Institutions, and the Great Divergence'. In: *Review of Economic Studies* 76.1, pp. 143–179.

Galor, Oded and Daniel Tsiddon (1997a). 'Technological Progress, Mobility, and Economic Growth'. In: *American Economic Review* 87.3, pp. 363–82.

– (1997b). 'The Distribution of Human Capital and Economic Growth'. In: *Journal of Economic Growth* 2.1, pp. 93–124.

Galor, Oded and Joseph Zeira (1993). 'Income Distribution and Macroeconomics'. In: *Review of Economic Studies* 60.1, pp. 35–52.

García-Peñalosa, Cecilia and Stephen J. Turnovsky (2011). 'Taxation and Income Distribution Dynamics in a Neoclassical Growth Model'. In: *Journal of Money, Credit and Banking* 43.8, pp. 1543–1577.

Glomm, Gerhard and B. Ravikumar (1998). 'Increasing returns, human capital, and the Kuznets curve'. In: *Journal of Development Economics* 55.2, pp. 353–367.

Goos, Maarten, Alan Manning and Anna Salomons (2014). 'Explaining Job Polarization: Routine-Biased Technological Change and Offshoring'. In: *American Economic Review* 104.8, pp. 2509–26.

Gradstein, Mark (2007). 'Inequality, democracy and the protection of property rights'. In: *Economic Journal* 117.516, pp. 252–269.

Greene, William H. (2003). *Econometric Analysis*. 5th. Upper Saddle River, NJ: Prentice Hall.

Greenwood, Jeremy and Boyan Jovanovic (1990). 'Financial Development, Growth, and the Distribution of Income'. In: *Journal of Political Economy* 98.5, pp. 1076–1107.

Grossman, Gene, Elhanan Helpman and Philipp Kircher (2013). 'Matching and Sorting in a Global Economy'. CEPR Discussion Papers No. 9680.

Grossman, Gene and Esteban Rossi-Hansberg (2008). 'Trading Tasks: A Simple Theory of Offshoring'. In: *American Economic Review* 98.5, pp. 1978–97.

– (2012). 'Task Trade Between Similar Countries'. In: *Econometrica* 80.2, pp. 593–629.

Grossmann, Volker (2008). 'Risky human capital investment, income distribution, and macroeconomic dynamics'. In: *Journal of Macroeconomics* 30.1, pp. 19–42.

Hanushek, Eric A., Charles Ka Yui Leung and Kuzey Yilmaz (2014). 'Borrowing Constraints, College Aid, and Intergenerational Mobility'. In: *Journal of Human Capital* 8.1, pp. 1–41.

Hartog, Joop (1979). *Personal Income Distribution: A Multicapability Theory*. Martinus Nijhoff Publishing.

Hausman, Jerry and William Taylor (1981). 'Panel Data and Unobservable Individual Effects'. In: *Econometrica* 49.6, pp. 1377–98.

Heckman, James (1976). 'The Common Structure of Statistical Models of Truncation, Sample Selection and Limited Dependent Variables and a Simple Estimator for Such Models'. In: *Annals of Economic and Social Measurement*. Vol. 5. NBER Chapters 4. National Bureau of Economic Research, Inc, pp. 475–492.

– (1979). 'Sample Selection Bias as a Specification Error'. In: *Econometrica* 47.1, pp. 153–61.

Heckman, James, John Eric Humphries and Gregory Veramendi (2016). *Returns to Education: The Causal Effects of Education on Earnings, Health and Smoking*. Working Paper 22291. National Bureau of Economic Research.

Heckman, James and Tim Kautz (2013). *Fostering and Measuring Skills: Interventions That Improve Character and Cognition*. NBER Working Papers 19656. National Bureau of Economic Research, Inc.

Herrendorf, Berthold, Richard Rogerson and Ákos Valentinyi (2014). 'Growth and Structural Transformation'. In: *Handbook of Economic Growth*. Vol. 2. Handbook of Economic Growth. Elsevier. Chap. 6, pp. 855–941.

Higgins, Matthew and Jeffrey G. Williamson (1999). *Explaining Inequality the World Round: Cohort Size, Kuznets Curves, and Openness*. NBER Working Papers 7224.

Holmes, Craig and Ken Mayhew (2015). *Have UK earnings distributions polarised?* INET Oxford Working Paper 2015-02.

Imbert, Clément (2013). 'Decomposing the Labor Market Earnings Inequality: The Public and Private Sectors in Vietnam, 1993–2006'. In: *World Bank Economic Review* 27.1, pp. 55–79.

International Federation of Robotics (2016). *IFR World Robotics 2016*.

Jacobsen, Joyce P., Melanie Khamis and Mutlu Yuksel (2015). 'Convergences in Men's and Women's Life Patterns: Lifetime Work, Lifetime Earnings, and Human Capital Investment'. In: *Research in Labor Economics*. Ed. by Solomon W. Polachek. Ed. by Konstantinos Tatsiramos. Ed. by Klaus F. Zimmermann. Vol. 41. Research in Labor Economics. Emerald Group Publishing Limited, pp. 1–33.

Jaremski, Matt, Elisabeth Perlman and Steven Schuster (2016). 'An Empirical History of the United States Postal Savings System'.

Jayadev, Arjun, Rahul Lahoti and Sanjay Reddy (2015). *The Middle Muddle: Conceptualizing and Measuring the Global Middle Class*. Courant Research Centre: Poverty, Equity and Growth - Discussion Papers 193.

Jenkins, Stephen (2004). 'Fitting functional forms to distributions, using ml'.

Kambourov, Gueorgui and Iourii Manovskii (2009). 'Occupational Specificity Of Human Capital'. In: *International Economic Review* 50.1, pp. 63–115.

Keane, Michael P and Kenneth I Wolpin (1997). 'The Career Decisions of Young Men'. In: *Journal of Political Economy* 105.3, pp. 473–522.

Khalifa, Sherif and Sherine El Hag (2010). 'Income Disparities, Economic Growth, And Development As A Threshold'. In: *Journal of Economic Development* 35.2, pp. 23–36.

Killingsworth, Mark and James Heckman (1987). 'Female labor supply: A survey'. In: *Handbook of Labor Economics*. Ed. by O. Ashenfelter and R. Layard. 1st ed. Vol. 1. Elsevier. Chap. 02, pp. 103–204.

Klump, Rainer, Peter McAdam and Alpo Willman (2007). 'Factor Substitution and Factor-Augmenting Technical Progress in the United States: A Normalized Supply-Side System Approach'. In: *The Review of Economics and Statistics* 89.1, pp. 183–192.

Klump, Rainer, Peter McAdam and Alpo Willman (2012). 'The Normalized CES Production Function: Theory And Empirics'. In: *Journal of Economic Surveys* 26.5, pp. 769–799.

Kremer, Michael (1993). 'Population Growth and Technological Change: One Million B.C. to 1990'. In: *The Quarterly Journal of Economics* 108.3, pp. 681–716.

Kremer, Michael and Daniel L Chen (2002). 'Income Distribution Dynamics with Endogenous Fertility'. In: *Journal of Economic Growth* 7.3, pp. 227–58.

Kripfganz, Sebastian and Claudia Schwarz (2015). *Estimation of linear dynamic panel data models with time-invariant regressors*. Tech. rep. 1838. European Central Bank.

Krusell, Per, Anthony A. Smith and Jr. (1998). 'Income and Wealth Heterogeneity in the Macroeconomy'. In: *Journal of Political Economy* 106.5, pp. 867–896.

Kuznets, Simon (1955). 'Economic Growth and Income Inequality'. In: *The American Economic Review* 45.1, pp. 1–28.

– (1963). 'Quantitative Aspects of the Economic Growth of Nations: VIII. Distribution of Income by Size'. In: *Economic Development and Cultural Change* 11.2, Part 2, pp. 1–80.

Lee, Jong-Wha and Hanol Lee (2016). 'Human capital in the long run'. In: *Journal of Development Economics* 122.C, pp. 147–169.

Lemieux, Thomas (1998). 'Estimating the Effects of Unions on Wage Inequality in a Panel Data Model with Comparative Advantage and Nonrandom Selection'. In: *Journal of Labor Economics* 16.2, pp. 261–91.

– (2006). 'The "Mincer Equation" Thirty Years After Schooling, Experience, and Earnings'. In: *Jacob Mincer A Pioneer of Modern Labor Economics*. Ed. by Shoshana Grossbard. Boston, MA: Springer US, pp. 127–145.

Lewis, W. Arthur (1954). 'Economic Development with Unlimited Supplies of Labour'. In: *The Manchester School* 22.2, pp. 139–191.

León-Ledesma, Miguel A., Peter McAdam and Alpo Willman (2010). 'Identifying the Elasticity of Substitution with Biased Technical Change'. In: *American Economic Review* 100.4, pp. 1330–57.

Li, Hongyi and Heng-fu Zou (1998). 'Income Inequality Is Not Harmful for Growth: Theory and Evidence'. In: *Review of Development Economics* 2.3, pp. 318–34.

Lipps, Oliver (2009). *Attrition of Households and Individuals in Panel Surveys*. SOEP-papers on Multidisciplinary Panel Data Research 164. DIW Berlin, The German Socio-Economic Panel (SOEP).

Lochner, Lance and Enrico Moretti (2004). 'The Effect of Education on Crime: Evidence from Prison Inmates, Arrests, and Self-Reports'. In: *American Economic Review* 94.1, pp. 155–189.

Lovenheim, Michael F. (2011). 'The Effect of Liquid Housing Wealth on College Enrollment'. In: *Journal of Labor Economics* 29.4, pp. 741–771.

Lucas, Robert Jr. (1988). 'On the mechanics of economic development'. In: *Journal of Monetary Economics* 22.1, pp. 3–42.

Maddison, Angus (2003). *The World Economy Historical Statistics: Historical Statistics.* Development Centre Studies. OECD Publishing.

Mazumder, Bhashkar (2005). 'Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data'. In: *The Review of Economics and Statistics* 87.2, pp. 235–255.

Mazzolari, Francesca and Giuseppe Ragusa (2013). 'Spillovers from High-Skill Consumption to Low-Skill Labor Markets'. In: *The Review of Economics and Statistics* 95.1, pp. 74–86.

Mihaela, Paunescu, Gagea Gabriela, Paunescu Catalin, Pitigoi Gabriel and Elisabeta Nicolae (2013). 'Relationship Between General Intelligence and Motor Skills Learning Specific to Combat Sports'. In: *Procedia - Social and Behavioral Sciences* 84. The 3rd World Conference on Psychology, Counseling and Guidance, WCPCG-2012, pp. 728 –732.

Milanovic, B. (2016). *Global inequality : a new approach for the age of globalization.* Cambridge, Massachusetts: Harvard University Press.

Milligan, Kevin, Enrico Moretti and Philip Oreopoulos (2004). 'Does education improve citizenship? Evidence from the United States and the United Kingdom'. In: *Journal of Public Economics* 88.9-10, pp. 1667–1695.

Mincer, Jacob (1974). *Schooling, Experience, and Earnings.* National Bureau of Economic Research, Inc.

Moran, T. (2005). 'Kuznets's Inverted U-Curve Hypothesis: The Rise, Demise, and Continued Relevance of a Socioeconomic Law'. In: *Sociological Forum* 20.2, pp. 209–244.

Morand, Oliver F. (1999). 'Endogenous Fertility, Income Distribution, and Growth'. In: *Journal of Economic Growth* 4.3, pp. 331–349.

Moreno-Galbis, Eva and Thepthida Sopraseuth (2014). 'Job polarization in aging economies'. In: *Labour Economics* 27.0, pp. 44–55.

Moretti, Enrico (2004). 'Human Capital Externalities in Cities'. In: *Handbook of Regional and Urban Economics* 4, pp. 2243–2291.

Mundlak, Yair (1978). 'On the Pooling of Time Series and Cross Section Data'. In: *Econometrica* 46.1, pp. 69–85.

Nellas, Viki and Elisabetta Olivieri (2012). 'The Change of Job Opportunities: The Role of Computerization and Institutions'. SSRN Working Paper Series.

Nisbett, Richard E, Joshua Aronson, Clancy Blair, William Dickens, James Flynn, Diane F Halpern and Eric Turkheimer (2012). 'Intelligence: new findings and theoretical developments.' In: *American psychologist* 67.2, p. 130.

OECD (2015). *The Labour Share in G20 Economies.* International Labour Organization, Organisation for Economic Co-operation and Development.

Oldenski, Lindsay (2014). 'Offshoring and the Polarization of the U.S. Labor Market'. In: *Industrial and Labor Relations Review* 67.2.5, pp. 734–761.

Oreopoulos, Philip and Ryan Dunn (2013). 'Information and College Access: Evidence from a Randomized Field Experiment*'. In: *The Scandinavian Journal of Economics* 115.1, pp. 3–26.

Papageorgiou, Chris and Nor Azam Abdul Razak (2009). *Inequality, Human Capital and Development: Making the Theory Face the Facts*. MPRA Paper. University Library of Munich, Germany.

Perotti, Roberto (1993). 'Political Equilibrium, Income Distribution, and Growth'. In: *The Review of Economic Studies* 60.4, pp. 755–776.

– (1996). 'Growth, Income Distribution, and Democracy: What the Data Say'. In: *Journal of Economic Growth* 1.2, pp. 149–87.

Persson, Torsten and Guido Tabellini (1994). 'Is Inequality Harmful for Growth?' In: *American Economic Review* 84.3, pp. 600–621.

Pesaran, M Hashem and Larry W Taylor (1999). 'Diagnostics for IV Regressions'. In: *Oxford Bulletin of Economics and Statistics* 61.2, pp. 255–81.

Piketty, Thomas (1997). 'The Dynamics of the Wealth Distribution and the Interest Rate with Credit Rationing'. In: *The Review of Economic Studies* 64.2, pp. 173–189.

– (2014). *Capital in the Twenty-First Century*. Harvard University Press.

Rauch, James E. (1993). 'Economic Development, Urban Underemployment, and Income Inequality'. In: *Canadian Journal of Economics* 26.4, pp. 901–18.

Roe, Mark J. and Jordan I. Siegel (2011). 'Political instability: Effects on financial development, roots in the severity of economic inequality'. In: *Journal of Comparative Economics* 39.3, pp. 279–309.

Romer, Paul M. (1986). 'Increasing Returns and Long-Run Growth'. In: *Journal of Political Economy* 94.5, pp. 1002–1037.

Rosen, Sherwin (1974). 'Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition'. In: *Journal of Political Economy* 82.1, pp. 34–55.

– (1987). 'The theory of equalizing differences'. In: *Handbook of Labor Economics*. Ed. by O. Ashenfelter and R. Layard. Handbook of Labor Economics. Elsevier. Chap. 12, pp. 641–692.

Roser, Max (2016). *Life Expectancy*. URL: https://ourworldindata.org/life-expectancy/.

Roy, A. D. (1951). 'Some Thoughts on the Distribution of Earnings'. In: *Oxford Economic Papers* 3.2, pp. 135–146.

Rubin, Donald B. (1976). 'Inference and missing data'. In: *Biometrika* 63.3, pp. 581–592.

Sattinger, Michael (1975). 'Comparative Advantage and the Distributions of Earnings and Abilities'. In: *Econometrica* 43.3, pp. 455–68.

– (1980). *Capital and the Distribution of Labor Earnings*. Elsevier, North-Holland.

Sherk, James (2013). *Productivity and Compensation: Growing Together*. URL: http://www.heritage.org/research/reports/2013/07/productivity-and-compensation-growing-together.

Smits-Engelsman, Bouwien and Elisabeth L. Hill (2012). 'The Relationship Between Motor Coordination and Intelligence Across the IQ Range'. In: *Pediatrics* 130.4, e950–e956.

Snijders, T.A.B. and R.J. Bosker (2011). *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling*. SAGE Publications.

Stoeger, H and A Ziegler (2010). 'How Fine Motor Skills Influence the Assessment of High Abilities and Underachievement in Math'. In: *Journal for the Education of the Gifted* 34.2, pp. 195–219.

Suri, Tavneet (2011). 'Selection and Comparative Advantage in Technology Adoption'. In: *Econometrica* 79.1, pp. 159–209.

Susskind, Daniel (2016). 'Technology and Employment: Tasks, Capabilities, and Tastes'. unpublished. PhD thesis. University of Oxford.

Taylor, Marcia, John Brice, Nick Buck and Elaine Prentice-Lane (2010). *British Household Panel Survey User Manual Volume A: Introduction, Technical Report and Appendices*. ISER Working Paper Series 2008-05. Colchester: University of Essex.

Uhrig, S.C. Noah (2008). *The nature and causes of attrition in the British Household Panel Study*. ISER Working Paper Series 2008-05. Institute for Social and Economic Research.

United States, National Monetary Commission and Sixty-First Congress (1910). *Notes on the Postal Savings-Bank Systems of the Leading Countries*. Washington: Government Printing Office.

Winters, John V. (2013). 'Human capital externalities and employment differences across metropolitan areas of the USA'. In: *Journal of Economic Geography* 13.5, pp. 799–822.

Wooldridge, Jeffrey (2010). *Econometric Analysis of Cross Section and Panel Data*. 2nd Ed. MIT Press Books. The MIT Press.

World Bank (2016). *World Development Report 2016: Digital Dividends*. Tech. rep. Washington, DC: World Bank. doi:10.1596/978-1-4648-0671-1. License: Creative Commons Attribution CC BY 3.0 IGO.

Yeaple, Stephen (2005). 'A simple model of firm heterogeneity, international trade, and wages'. In: *Journal of International Economics* 65.1, pp. 1–20.

Young, Andrew (2013). 'U.S. Elasticities of Substitution and Factor-Augmentation at the Industry Level'. In: *Macroeconomic Dynamics* 17.04, pp. 861–897.

Appendices

Appendix A: Macroeconomic identities and Balance of Payments

This appendix presents the complete macroeconomic accounts of the model, together with an analysis of external accounts.

A.1 Macroeconomic identities

Gross domestic product and gross national products are respectively (recall that only old agents consume):

$$GDP_t = Y_t^s + Y_t^n = C_t^o + I_t$$

$$GNP_t = C_t^o + I_t - rK_t^{*F}$$

The last term (rK_t^{*F}) is equal to the net factor income payments (only capital is internationally mobile). To obtain the per capita expressions these need to be divided by $2N$.

Total savings are equal to investment:

$$S_t = GDP_t - C_t^o = I_t$$

But total savings include national and external savings. National savings are:

$$GNP_t - C_t^o = I_t - rK_t^{*F}$$

Which means that external savings are rK_t^{*F}. This is exactly the value of the current account, as expected.[1]

[1]The current account includes the trade balance (zero in this economy) plus transfers and net income paid to foreign factors. Only the latter is non-zero in this economy.

A.2 External balance

Under given parametrisations, the model always has a steady state (e.g. if **Condition 1** holds). Because there is nothing in the model that ensures equal demand and supply of capital, the steady state will very likely imply either import or export of capital, forever. Is a permanent imbalance in the capital market sustainable? The answer lies in the balance of payments (BoP henceforth). The BoP is composed of the sum of the current account and the capital account. The former accounts for the net factor income payments – in particular the net return to capital, whereas the latter reflects the net entrance of capital. As shown here, given a domestic capital market disequilibrium, the BoP will be always unbalanced, but with a finite present value. Therefore, it is sustainable.

To see this, consider the example of an economy at the steady state but with an excess demand for capital. At the beginning of the period, firms will import K^F units of capital, an amount reflected positively in the financial account. At the end of the period, after firms have used capital for production, they will give back the borrowed capital K^F, reflected negatively in the financial account. They also return an interest on this borrowed capital, equivalent to rK^F. This is reflected negatively in the current account. In result, the capital account is balanced whereas the current account deficit is equal to rK^F. The BoP deficit is therefore rK^F. This process repeats every period. To evaluate sustainability (defined as a BoP that is not explosive as $t \to \infty$), let us calculate the present value of all current account deficits. Since this is a perpetuity, the following formula applies:

$$PV(BoP) = \frac{rK^F}{r} = K^F < \infty$$

This is, the present value of the BoP is finite. This is true because firms are paying in interest the total amount of capital they borrow every period. This is equivalent to what is finance is called “equity financing”: issuing equity to buy assets (here, capital).

If the economy is exporting capital, there exist a permanent positive current account due to the interest paid to capital owned by local agents. This means higher agent consumption and bequests. The analysis in the main document only refers to output, and hence it is independent of the ownership of firms. Naturally, if some of the local firms are owned by foreign agents – through an inflow of capital, part of total output payment drains abroad to pay capital owners, lowering the welfare of the economy (measured in terms of consumption and bequests. In other words, the foreign sector and the external accounts are crucial for population's welfare, but not for the equilibrium of the model.

Finally, if the economy starts outside the steady state, the result holds but with a finite discrepancy in the present value of the BoP. During the transition period the economy will be importing (exporting) variable amounts of capital and paying (receiving) net interest for them. Once it reaches the steady state, it will be paying (receiving) a fixed amount, as above. The proof is trivial. At any time outside the equilibrium, the BoP is rK_t^F. This can always be decomposed as $rK^F + \epsilon_t$. If the economy starts in $t = 0$ and reaches the steady state in $t = T$, then the present value of the BoP is:

$$PV(BoP) = \sum_{t=0}^{T-1} \frac{\epsilon_t}{(1+r)^t} + K^F$$

where the second term is the perpetuity, as above. Assuming that the model always reaches a steady state (as in our standard Kuznets story), the first term is finite. Therefore, sustainability holds.

Appendix B: Total equilibrium private aggregate demand for skilled labour

This appendix studies the functional form of the “total” equilibrium aggregate demand for skilled labour. This demand is:

$$w_t^s = A^s \left(L_t^s\right)^{\frac{\psi}{1-\theta}} (1-\theta) \left(\frac{A^s \theta}{r}\right)^{\frac{\theta}{1-\theta}} \tag{B.1}$$

Define:

$$Z = A^s (1-\theta) \left(\frac{A^s \theta}{r}\right)^{\frac{\theta}{1-\theta}} > 0 \tag{B.2}$$

Then the first derivative of this demand is:

$$\frac{\partial w_t^s}{\partial L_t^s} = \left(\frac{\psi}{1-\theta}\right) L_t^{s\left(\frac{\psi-(1-\theta)}{1-\theta}\right)} Z \tag{B.3}$$

This is positive for $L_t^s > 0$, and for any positive ψ. Thus, regardless of the shape of the “partial” demand, the “total” demand is always upward slopping. As always, $\psi = 0$ turns this derivative equal to zero, meaning the demand is perfectly elastic, as in GZ93.

The second derivative is:

$$\frac{\partial^2 w_t^s}{\partial L_t^{s2}} = \left(\frac{\psi-(1-\theta)}{1-\theta}\right) \left(\frac{\psi}{1-\theta}\right) L_t^{s\left(\frac{\psi-2(1-\theta)}{1-\theta}\right)} Z \tag{B.4}$$

The sign of this expression depends on the term $\psi - (1-\theta)$, in the first parenthesis. If positive, the demand is strictly convex; if negative, it is strictly concave; if zero, it is a straight line. This property comes from equation (1.8), which states how fast the capital *per skilled worker* grows (notice the exponent of skilled labour is equal to the one in equation B.1, namely $\frac{\psi}{1-\theta}$. If $k_t^{D^*}$ more than doubles when L_t^s doubles (i.e. $\psi > (1-\theta)$), aggregate capital $K_t^{D^*}$ grows much faster (at an exponent larger than 2) than L_t^s, producing a demand that accelerates with skilled labour; the second derivative is positive. On the contrary, if $k_t^{D^*}$ grows by less than the growth of L_t^s (i.e. $\psi < (1-\theta)$), aggregate capital $K_t^{D^*}$ grows faster than labour but not too fast (at an exponent between 1 and 2), producing a demand that is increasing but at a decelerating rate; the second derivative is negative. Intuitively, the productivity of skilled workers (and their wage) comes from two sources: (i) the aggregate productivity level, and (ii) the level of capital *per skilled worker* (see equation 1.9). The former increases productivity at rate ψ with respect to increases in L^s (from equation 1.4), whereas the latter increases productivity at rate θ (the elasticity of capital in the Cobb-Douglas function). Thus, the productivity of the skilled wage is increasing at an accelerating rate whenever the sum of these two effects is larger than one; this is, when $\psi + \theta > 1$, equivalent to $\psi - (1 - \theta) > 0$.

Appendix C: Distributions of ability

The endowment of ability for each worker on each task is given by equation (2.2). This is chosen because it "gives origin" to a log normal distribution with parameters μ_j and σ_j^2, corresponding to the mean and variance of the adjacent normal distribution. This is true in the sense that a histogram of the ability of all workers resembles that of a log-normal distribution, but not in a stochastic sense. The model presented here is fully deterministic.

To see this relation, start from equation (2.2). The inverse of this function is:

$$i(\eta_j) = \frac{1}{2}\left[1 + \text{erf}\left(\frac{\ln(\eta_j) - \mu_j}{\sigma_j\sqrt{2}}\right)\right] \tag{C.1}$$

This is exactly the cumulative distribution function (CDF) of a log-normal distribution, where $i \in [0, 1]$. Now, consider the following method:

Inverse Transform Method

For any continuous random variable X, its CDF $F(x) = Y$ follows a standard uniform distribution, i.e. $Y \sim U[0, 1]$. As a consequence, $F^{-1}(y)$ follows the distribution of X.

Recall that workers are defined as a continuous mass of size 1 over $[0, 1]$, This is, i – and $i(\eta_j)$ – $\in [0, 1]$. Additionally, $\eta_j(i)$ is continuous, monotonic, and increasing over i (see Panel (a) in Figure 2.3), which means $i(\eta_j)$ is also continuous, monotonic, and increasing over η_j. These two properties combined imply $i(\eta_j)$ is equivalent to the **proportion** of workers with an ability level equal or lower than η_j. For instance, $i(1.8) = 0.5$ means 50% of workers have an ability equal or lower than 1.8 at this task. These properties define a CDF. So, shifting to the Method's notation, $i(\eta_j)$ is equivalent to Y, the CDF of X. Now, we already found that this CDF corresponds to a log-normal distribution. Then X follows a log-normal distribution. Applying the Method, it follows that the inverse of Y – $\eta_j(i)$ in our case – follows the distribution of X, i.e. a log-normal distribution.

Appendix D: Workers preferences

Labour allocation depends on the relative order of i_1, i_2, and i_3. The main text presented the only possible scenario consistent with the labour market. Regarding the other five cases, one is consistent with workers preferences but it is inconsistent with the demand for tasks. Panel (a) in Figure A1 presents this case, where workers supply labour for two tasks only – m and a. Clearly, with a CES production function with imperfect substitution ($\sigma_Y < \infty$), the task r is essential for output. Its demand is then positive. In consequence, this scenario cannot occur in equilibrium.[1] The remaining four cases cannot occur because they are inconsistent with workers preferences. Panels (b) and (c) in Figure A1 show two of these scenarios, where preferences are intransitive for a certain group of workers.

[1]If no labour is supplied for this task, the relative wage for ω_r will rise. In turn, i_1 will decrease and i_2 will increase until the inequality is reversed – which is the scenario that does occur in equilibrium.

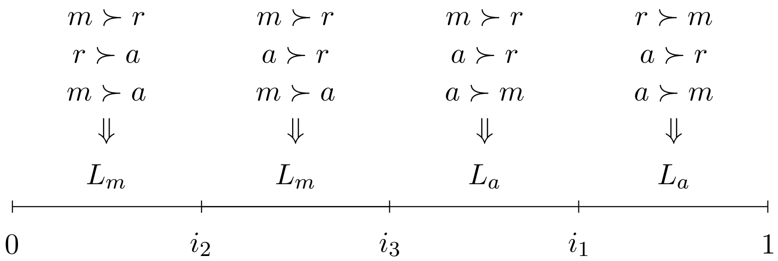

(a) Labour allocation if $i_2 < i_3 < i_1$

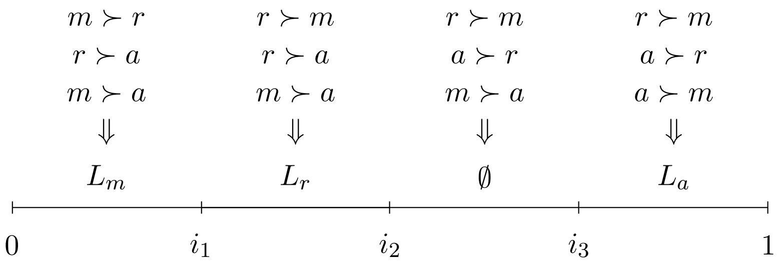

(b) Labour allocation if $i_1 < i_2 < i_3$

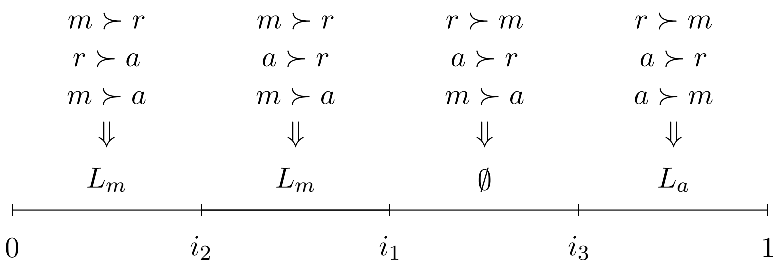

(c) Labour allocation if $i_2 < i_1 < i_3$

Figure A1: Labour market allocation not possible in equilibrium

Appendix E: General Solution

E.1 Three tasks

Consider a set of continuous, monotonic and increasing functions $\{\eta_m(i), \eta_r(i), \eta_a(i)\}$. Furthermore, assume the following:

Strict monotonicity

$$\frac{\partial\left(\frac{\eta_r(i)}{\eta_m(i)}\right)}{\partial i} > 0 \quad \forall i \quad \text{(A.1)} \qquad\qquad \frac{\partial\left(\frac{\eta_a(i)}{\eta_r(i)}\right)}{\partial i} > 0 \quad \forall i \quad \text{(A.2)}$$

These assumptions mean that the pairwise comparison of ability across workers is strictly monotonic over all i. Figure A1 shows these assumptions graphically. There are two key implications of A.1 and A.2, visible in the graphs. First, the strict monotonicity of these rations (given by the strict inequality sign) implies that for each pairwise comparison of abilities there is only one indifferent worker among them. Thus, there are three cut-off points along i (i_1, i_2, i_3). Second, the specific sign chosen ($>$) means that the only possible ordering among them in equilibrium is $i_1 < i_3 < i_2$ (not shown). This means that the ordering of workers to tasks along i is $m - r - a$, resulting in the following labour allocation:

$$L_m = i_1 \qquad\qquad L_r = i_2 - i_1 \qquad\qquad L_a = 1 - i_2$$

where $L_m + L_r + L_a = 1$.

Solution

For this generic case, the maximisation problem is:

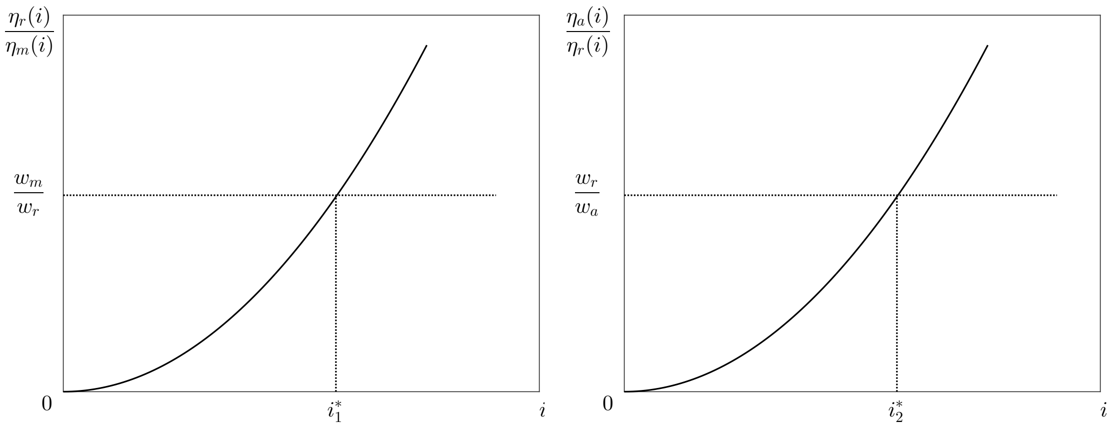

(a) Assumption A.1 (b) Assumption A.2

Figure A1: Uniqueness of equilibrium and task ordering assured by assumptions A.1 and A.2

Maximisation problem

$$\max_{\{i_1, i_2\}} \quad Y$$

subject to:

$$\begin{aligned} Y &= \left(\alpha_m T_m^{\phi} + \alpha_r T_r^{\phi} + \alpha_a T_a^{\phi}\right)^{\frac{1}{\phi}} \\ T_a &= g(i_2) \\ T_r &= g(i_1, i_2) \\ T_m &= g(i_1) \end{aligned}$$

The respective first order conditions are:

$$\begin{aligned} \left(\frac{T_r}{T_m}\right)^{1-\phi} &= \left(\frac{\alpha_r}{\alpha_m}\right)\left(-\frac{\frac{\partial T_r}{\partial i_1}}{\frac{\partial T_m}{\partial i_1}}\right) \\ \left(\frac{T_a}{T_r}\right)^{1-\phi} &= \left(\frac{\alpha_a}{\alpha_r}\right)\left(-\frac{\frac{\partial T_a}{\partial i_2}}{\frac{\partial T_r}{\partial i_2}}\right) \end{aligned}$$

Since $T_j \equiv \int_a^b \eta_j(i)di$, using the Leibniz' rule (about differentiation of an integral) these FOCs can be further expressed as (* denotes optimal):[1]

[1]Given the definition of T_j, the Leibniz's rule implies that $\frac{\partial T_j}{\partial b} = \eta_j(b)$ and $\frac{\partial T_j}{\partial a} = -\eta_j(a)$.

2x2 System

$$\left(\frac{T_r}{T_m}\right)^{1-\phi} = \left(\frac{\alpha_r}{\alpha_m}\right)\left(\frac{\eta_r(i_1^*)}{\eta_m(i_1^*)}\right)$$
$$\left(\frac{T_a}{T_r}\right)^{1-\phi} = \left(\frac{\alpha_a}{\alpha_r}\right)\left(\frac{\eta_a(i_2^*)}{\eta_r(i_2^*)}\right)$$

The solution of the model comes from i_1^* and i_2^*, obtained from the above system of equations.

E.2 N tasks

Consider a set of continuous, monotonic and increasing functions $\{\eta_{T_1}(i), \eta_{T_2}(i), ..., \eta_{T_N}(i)\}$, where η_{T_j} denotes ability of worker i in task T_j. Furthermore, assume the following (ordering does not deter the generality of the setup):

Strict monotonicity

$$\frac{\partial\left(\frac{\eta_{T_2}(i)}{\eta_{T_1}(i)}\right)}{\partial i} > 0 \quad \forall i \quad , \quad \frac{\partial\left(\frac{\eta_{T_3}(i)}{\eta_{T_2}(i)}\right)}{\partial i} > 0 \quad \forall i \quad , \quad \dots \quad , \quad \frac{\partial\left(\frac{\eta_{T_N}(i)}{\eta_{T_{N-1}}(i)}\right)}{\partial i} > 0 \quad \forall i$$

Given these assumptions it is possible to show that there exists $\binom{N}{2}$ different points over i such that at each of these points two workers are indifferent among a given pair of tasks. Yet, only N points are relevant for determining the allocation of workers into tasks. This is, the i dimension is divided into N intervals, where at each of them the corresponding workers allocate to one task in following fashion:

$$L_1 = i_1 \qquad L_2 = i_2 - i_1 \qquad \dots \qquad L_N = 1 - i_N$$

where $\sum_{k=1}^{N} L_k = 1$.

Solution

The maximisation problem is the following:

Maximisation problem

$$\max_{\{i_1,\dots,i_N\}} \quad Y$$

subject to:

$$\begin{aligned} Y &= \left(\alpha_1 T_1^{\phi} + \cdots + \alpha_N T_N^{\phi}\right)^{\frac{1}{\phi}} \\ T_N &= g(i_N) \\ T_{N-1} &= g(i_{N-1}, i_N) \\ &\vdots \\ T_2 &= g(i_1, i_2) \\ T_1 &= g(i_1) \end{aligned}$$

The solution to this problem comes from the following $N \times N$ system of equations:

$N \times N$ System

$$\begin{aligned} \left(\frac{T_2}{T_1}\right)^{1-\phi} &= \left(\frac{\alpha_2}{\alpha_1}\right)\left(\frac{\eta_2(i_1^*)}{\eta_1(i_1^*)}\right) \\ &\vdots \\ \left(\frac{T_N}{T_{N-1}}\right)^{1-\phi} &= \left(\frac{\alpha_N}{\alpha_{N-1}}\right)\left(\frac{\eta_N(i_N^*)}{\eta_{N-1}(i_N^*)}\right) \end{aligned}$$

Appendix F: Characterisation of the implicit function for i_1^* and i_2^*

Consider the two-equation system defined by (2.18) and (2.19). Redefine each equation as follows:

$$F^1(i_1, i_2, \alpha_j, \mu_j, \sigma_j, \phi) \equiv \left(\frac{T_r}{T_m}\right)^{1-\phi} - \left(\frac{\alpha_r}{\alpha_m}\right) exp\left[(\mu_r - \mu_m) + (\sigma_r - \sigma_m)\sqrt{2}\,\mathrm{erf}^{-1}\left(2i_1^* - 1\right)\right] = 0$$

$$F^2(i_1, i_2, \alpha_j, \mu_j, \sigma_j, \phi) \equiv \left(\frac{T_r}{T_a}\right)^{1-\phi} - \left(\frac{\alpha_r}{\alpha_a}\right) exp\left[(\mu_r - \mu_a) + (\sigma_r - \sigma_a)\sqrt{2}\,\mathrm{erf}^{-1}\left(2i_2^* - 1\right)\right] = 0$$

where the endogenous variables are i_1 and i_2, and the exogenous parameters are α_j, μ_j, σ_j, and ϕ. The Jacobian of this system is defined as:

$$J = \begin{bmatrix} \frac{\partial F^1}{\partial i_1} & \frac{\partial F^1}{\partial i_2} \\ \frac{\partial F^2}{\partial i_1} & \frac{\partial F^2}{\partial i_2} \end{bmatrix}$$

Now, we can introduce the Implicit Function Theorem:

Theorem 1. *Assume $F^1(\cdot)$ and $F^2(\cdot)$ are real-valued functions defined on a domain Θ, and continuously differentiable over $\Theta^1 \subset \Theta$. For a given parametrisation $(\alpha_j, \mu_j, \sigma_j, \phi)$, consider point $(i_1^*, i_2^*) \in \Theta^1$, such that $F^1(i_1^*, i_2^*) = 0$ and $F^2(i_1^*, i_2^*) = 0$. Furthermore, assume the determinant of the Jacobian of these functions with respect to i_1 and i_2, evaluated at (i_1^*, i_2^*) is non-zero. Then, there exists a neighbourhood around $\alpha_j, \mu_j, \sigma_j, \phi$ such that the real-valued functions below exist:*

- $i_1^* = f_1(\alpha_j, \mu_j, \sigma_j, \phi)$
- $i_2^* = f_2(\alpha_j, \mu_j, \sigma_j, \phi)$

Additionally, the partial derivatives of these functions with respect to any given parameter are also defined, and can be found from the Cramer's Rule as follows:

$$\frac{\partial f_1}{\partial \alpha_j} = -\frac{\begin{vmatrix} \frac{\partial F^1}{\partial \alpha_j} & \frac{\partial F^1}{\partial i_2} \\ \frac{\partial F^2}{\partial \alpha_j} & \frac{\partial F^2}{\partial i_2} \end{vmatrix}}{\begin{vmatrix} \frac{\partial F^1}{\partial i_1} & \frac{\partial F^1}{\partial i_2} \\ \frac{\partial F^2}{\partial i_1} & \frac{\partial F^2}{\partial i_2} \end{vmatrix}} \qquad \frac{\partial f_2}{\partial \alpha_j} = -\frac{\begin{vmatrix} \frac{\partial F^1}{\partial i_1} & \frac{\partial F^1}{\partial \alpha_j} \\ \frac{\partial F^2}{\partial i_1} & \frac{\partial F^2}{\partial \alpha_j} \end{vmatrix}}{\begin{vmatrix} \frac{\partial F^1}{\partial i_1} & \frac{\partial F^1}{\partial i_2} \\ \frac{\partial F^2}{\partial i_1} & \frac{\partial F^2}{\partial i_2} \end{vmatrix}}$$

and similarly for the rest of parameters.

Let us apply this theorem. First, a quick look at the implicit functions F^1 and F^2 reveal that they are continuous over all the parameter space.[1] The only possible discontinuities would exist if $\alpha_j = 0$ or $T_j = 0$, for any j. These however are ruled out from the parameter space (the latter because of $\phi < 1$). Additionally, as shown below, the determinant of the Jacobian is never zero, provided $\phi < 1$. These two imply the above theorem applies to the whole parameter space.

Second, let us calculate the Jacobian matrix. The derivation of each of these elements comes from a rather mechanical application of the chain rule, and thus are not shown. The results are:

$$\frac{\partial F^1}{\partial i_1} = -\left(\frac{T_r}{T_m}\right)^{1-\phi}\left[(1-\phi)\left[\frac{\eta_r}{T_r}+\frac{\eta_m}{T_m}\right] + (\sigma_r - \sigma_m)\sqrt{(2\pi)}exp\left(\mathrm{erf}^{-1}\left(2i_1^* - 1\right)^2\right)\right]$$

$$\frac{\partial F^1}{\partial i_2} = (1-\phi)\left(\frac{T_r}{T_m}\right)^{1-\phi}\frac{\eta_r}{T_r}$$

$$\frac{\partial F^2}{\partial i_1} = -(1-\phi)\left(\frac{T_r}{T_a}\right)^{1-\phi}\frac{\eta_r}{T_r}$$

$$\frac{\partial F^2}{\partial i_2} = \left(\frac{T_r}{T_a}\right)^{1-\phi}\left[(1-\phi)\left[\frac{\eta_r}{T_r}+\frac{\eta_a}{T_a}\right] + (\sigma_a - \sigma_r)\sqrt{(2\pi)}exp\left(\mathrm{erf}^{-1}\left(2i_2^* - 1\right)^2\right)\right]$$

Now, we can calculate the determinant of J, which corresponds to $\frac{\partial F^1}{\partial i_1} \times \frac{\partial F^2}{\partial i_2} - \frac{\partial F^1}{\partial i_2} \times$

[1] This space is defined by $-\infty < \phi < 1$, $\alpha_j > 0$, $\mu_j > 0$ and $\sigma_j > 0$, $\forall j$, together with $\sigma_m < \sigma_r < \sigma_a$ (**Assumption 1**).

$\dfrac{\partial F^2}{\partial i_1}$. After some rearrangements, this is:

$$|J| = -\left(\frac{T_r}{T_m}\right)^{1-\phi}\left(\frac{T_r}{T_a}\right)^{1-\phi}\Bigg[(1-\phi)^2\left(\frac{\eta_m}{T_m}\frac{\eta_r}{T_r}+\frac{\eta_a}{T_a}\frac{\eta_r}{T_r}+\frac{\eta_m}{T_m}\frac{\eta_a}{T_a}\right)$$
$$+(1-\phi)\left(\frac{\eta_m}{T_m}+\frac{\eta_r}{T_r}\right)(\sigma_a-\sigma_r)\sqrt{2\pi}exp\left(\mathrm{erf}^{-1}\left(2i_2^*-1\right)^2\right)$$
$$+(1-\phi)\left(\frac{\eta_a}{T_a}+\frac{\eta_r}{T_r}\right)(\sigma_r-\sigma_m)\sqrt{2\pi}exp\left(\mathrm{erf}^{-1}\left(2i_1^*-1\right)^2\right)$$
$$+2\pi(\sigma_a-\sigma_r)^2(\sigma_r-\sigma_m)^2\left[exp\left(\mathrm{erf}^{-1}\left(2i_1^*-1\right)^2\right)exp\left(\mathrm{erf}^{-1}\left(2i_2^*-1\right)^2\right)\right]^2\Bigg]$$

All the terms above are always positive over the parameter space (notice the role of **Assumption 1**). As such, the determinant of the Jacobian is never zero, and always negative (given by the minus outside the expression).

To calculate the sign of any derivative, we will use the Cramer's Rule, as outlined in the theorem above. Notice that there is a minus outside the fraction. Given the negative sign of the denominator just found, the sign of the final partial derivative is equal to the sign of the respective nominator's determinant. Here I show some examples.

- α_r

According to the Implicit Function Theorem, the derivative of i_1^* with respect to this parameter is given by:

$$\frac{\partial f_1}{\partial \alpha_r} = -\frac{\begin{vmatrix}\frac{\partial F^1}{\partial \alpha_r} & \frac{\partial F^1}{\partial i_2}\\ \frac{\partial F^2}{\partial \alpha_r} & \frac{\partial F^2}{\partial i_2}\end{vmatrix}}{\begin{vmatrix}\frac{\partial F^1}{\partial i_1} & \frac{\partial F^1}{\partial i_2}\\ \frac{\partial F^2}{\partial i_1} & \frac{\partial F^2}{\partial i_2}\end{vmatrix}} \tag{F.1}$$

After a bit of rearranging, the derivatives of F^1 and F^2 with respect to α_r are:

$$\frac{\partial F^1}{\partial \alpha_r} = -\left(\frac{T_r}{T_m}\right)^{1-\phi}\frac{1}{\alpha_r}$$

$$\frac{\partial F^2}{\partial \alpha_r} = -\left(\frac{T_r}{T_a}\right)^{1-\phi}\frac{1}{\alpha_r}$$

Combining this with the already calculated elements of the Jacobian, the determinant of the numerator in (F.1) is:

$$-\frac{1}{\alpha_r}\left(\frac{T_r}{T_m}\right)^{1-\phi}\left(\frac{T_r}{T_a}\right)^{1-\phi}\left[(1-\phi)\frac{\eta_a}{T_a}+(\sigma_a-\sigma_r)\sqrt{2\pi}\,exp\left(\text{erf}^{-1}\left(2i_2^*-1\right)^2\right)\right]$$

Given the parameter space being used, this expression is **always** negative.

The derivative of i_2^* with respect to α_r comes from an equivalent procedure. The determinant of the respective numerator is:

$$\frac{1}{\alpha_r}\left(\frac{T_r}{T_m}\right)^{1-\phi}\left(\frac{T_r}{T_a}\right)^{1-\phi}\left[(1-\phi)\frac{\eta_m}{T_m}+(\sigma_r-\sigma_m)\sqrt{2\pi}\,exp\left(\text{erf}^{-1}\left(2i_1^*-1\right)^2\right)\right]$$

which is **always** positive. In consequence, an increase in α_r leads to an unambiguous fall in L_m and L_a. Finally, the derivatives for α_m and α_a are calculated similarly, and not shown.

- μ_r

In this case, the derivative of i_1^* with respect to μ_r is:

$$\frac{\partial f_1}{\partial \mu_r} = -\frac{\begin{vmatrix}\frac{\partial F^1}{\partial \mu_r} & \frac{\partial F^1}{\partial i_2}\\ \frac{\partial F^2}{\partial \mu_r} & \frac{\partial F^2}{\partial i_2}\end{vmatrix}}{\begin{vmatrix}\frac{\partial F^1}{\partial i_1} & \frac{\partial F^1}{\partial i_2}\\ \frac{\partial F^2}{\partial i_1} & \frac{\partial F^2}{\partial i_2}\end{vmatrix}} \tag{F.2}$$

The derivatives of F^1 and F^2 with respect to μ_r are:

$$\frac{\partial F^1}{\partial \alpha_r} = -\phi\left(\frac{T_r}{T_m}\right)^{1-\phi}$$

$$\frac{\partial F^2}{\partial \alpha_r} = -\phi\left(\frac{T_r}{T_a}\right)^{1-\phi}$$

After some rearranging, the determinant of the numerator in (F.2) is:

$$-\phi\left(\frac{T_r}{T_m}\right)^{1-\phi}\left(\frac{T_r}{T_a}\right)^{1-\phi}\left[(1-\phi)\frac{\eta_a}{T_a}+(\sigma_a-\sigma_r)\sqrt{2\pi}\,exp\left(\text{erf}^{-1}\left(2i_2^*-1\right)^2\right)\right]$$

In the parameters space under consideration the terms in brackets are always positive. As such, the sing of this derivative depends entirely on ϕ, as Table 2.2 shows.

The derivative of i_2^* with respect to μ_r comes from a similar procedure. The determinant of the respective numerator is:

$$-\phi\left(\frac{T_r}{T_a}\right)^{1-\phi}\left(\frac{T_r}{T_m}\right)^{1-\phi}\left[(1-\phi)\frac{\eta_m}{T_m}+(\sigma_a-\sigma_r)\sqrt{2\pi}\,exp\left(\operatorname{erf}^{-1}(2i_2^*-1)^2\right)\right]$$

which again depends solely on the sign of ϕ. The derivatives for μ_m and μ_a are also simple to calculate, and not shown.

- σ_r

The derivatives of F^1 and F^2 with respect to σ_r are:

$$\frac{\partial F^1}{\partial \sigma_r}=\sqrt{2}\left[(1-\phi)-\operatorname{erf}^{-1}(2i_1^*-1)\right]\left(\frac{T_r}{T_m}\right)^{1-\phi}$$

$$\frac{\partial F^2}{\partial \sigma_r}=\sqrt{2}\left[(1-\phi)-\operatorname{erf}^{-1}(2i_2^*-1)\right]\left(\frac{T_r}{T_a}\right)^{1-\phi}$$

The determinant of the numerator of the derivative of i_1^* with respect to σ_r is:

$$\sqrt{2}\left(\frac{T_r}{T_m}\right)^{1-\phi}\left(\frac{T_r}{T_a}\right)^{1-\phi}\Bigg[\left[(1-\phi)-\operatorname{erf}^{-1}(2i_1^*-1)\right]\Big((1-\phi)\frac{\eta_a}{T_a} + (\sigma_a-\sigma_r)\sqrt{2\pi}\,exp\left(\operatorname{erf}^{-1}(2i_2^*-1)^2\right)\Big) + \left((1-\phi)\frac{\eta_r}{T_r}\left[\operatorname{erf}^{-1}(2i_2^*-1)-\operatorname{erf}^{-1}(2i_1^*-1)\right]\right)\Bigg]$$

The sign of this expression depends on the sing of the long squared bracket. The second component is always positive, because $i_2^* > i_1^*$, and the inverse error function is strictly monotonic and increasing. The first component has to parts. The one to the right is always positive, whereas the one to the left might be positive or negative, depending on the **level of** i_1^*. If this is very close to one, this expression could become negative, because $\operatorname{erf}^{-1}(1)=\infty$. Yet, this is the case of an economy with a very large manual occupations share, unlikely to represent any advanced economy – our main interest.

The case of the derivative of i_2^* with respect to σ_r is symmetrical with the above, which means that now the last term in the above expression is $\left[\operatorname{erf}^{-1}(2i_1^*-1)-\operatorname{erf}^{-1}(2i_2^*-1)\right]$, which is always negative. Singe the other term is the multiplication of positive and negative components, the likelihood of a change in sign of derivatives is now greater than before. That is the case shown in Table 2.2. Again, this is more likely to occur if i_1^* and i_2^* are close to zero (i.e. large L_a), close to one (i.e. large L_m), or very far apart from each other (i.e. large L_r). This are certainly not the most realistic cases.

The derivatives for σ_m and σ_a follow the same method, and not shown.

- ϕ

The derivatives of F^1 and F^2 with respect to ϕ are:

$$\frac{\partial F^1}{\partial \phi} = -\left(\frac{T_r}{T_m}\right)^{1-\phi} \ln\left(\frac{T_r}{T_m}\right)$$

$$\frac{\partial F^2}{\partial \phi} = -\left(\frac{T_r}{T_a}\right)^{1-\phi} \ln\left(\frac{T_r}{T_a}\right)$$

The determinant of the numerator corresponding to the derivative of i_1^* with respect to ϕ is:

$$\left(\frac{T_r}{T_m}\right)^{1-\phi}\left(\frac{T_r}{T_a}\right)^{1-\phi}\left[(1-\phi)\frac{\eta_r}{T_r}\ln\left(\frac{T_m}{T_a}\right) + \ln\left(\frac{T_r}{T_m}\right)\left[(1-\phi)\frac{\eta_a}{T_a} + (\sigma_a - \sigma_r)\sqrt{2\pi}\, exp\left(\text{erf}^{-1}\left(2i_2^* - 1\right)^2\right)\right]\right]$$

The sign of this expression depends on the intensity of use of different tasks in production. For example, the derivative under study is **positive** (i.e. L_m increases with ϕ) if $T_m > T_a$ and $T_m > T_r$. This is, employment in manual tasks expands when its use is the largest among inputs. This is intuitive. If there is greater substitution between tasks, firms can focus on using those tasks which are more productive (T_m). The converse is also true. Namely, if T_m is the lowest of tasks, the increase in ϕ reduces L_m.

A symmetric result occurs when looking at the derivative of i_2^* with respect to ϕ:

$$\left(\frac{T_r}{T_m}\right)^{1-\phi}\left(\frac{T_r}{T_a}\right)^{1-\phi}\left[(1-\phi)\frac{\eta_r}{T_r}\ln\left(\frac{T_m}{T_a}\right) + \ln\left(\frac{T_r}{T_a}\right)\left[(1-\phi)\frac{\eta_m}{T_m} + (\sigma_r - \sigma_m)\sqrt{2\pi}\, exp\left(\text{erf}^{-1}\left(2i_1^* - 1\right)^2\right)\right]\right]$$

This term is **negative** (i.e. L_a increases) if T_a is larger than the other inputs. Again, greater substitution redirects the input mix in the production of the final good towards most productive tasks. If T_a is the lowest among inputs, greater substitution shift employment away of L_a.

These and the previous results are of course linked in such a way that any change in L_r is ambiguous. For example, if $T_a < T_r < T_m$, an increase in ϕ expands L_m and reduces L_a. Just as for the case of other parameters, this "downward" mobility denies the identification of changes to L_r. Alternatively, if $T_r < T_a < T_m$, an increase in ϕ expands L_m, without providing unambiguous information about the other variables. In fact, a look at the equations above show that the final change in shares depend on the relative value of all the other parameters, including ϕ and the equilibrium values. For example, in the case of $T_r < T_a < T_m$, an increase in ϕ might lead to polarisation if i_2^* falls, which is more likely the larger i_1^* or the term $(\sigma_r - \sigma_m)$ are. It is not clear how much intuition there is behind these results.

Finally, the reason why a given T_j is greater than the rest hinges on three factors: (i) the magnitude of $\mu_j - \mu_{-j}$; (ii) the magnitude of $\sigma_j - \sigma_{-j}$; and (iii) the level of the equilibria i_1^* and i_2^* (which define the magnitude of T_j). It is not possible to generalise from these, except for the intuitive result that the more able workers are in a particular task, the more its use expands in production due to an increase in the degree of substitution.[2]

[2]Still, notice that the impact of μ_j and σ_j on L_j depends on the value of ϕ (as this appendix and Table 2.2 shows), making the analysis more complex.

Appendix G: Selection bias

The ultimate goal of the exercise carried out in Chapter 3 is to identify the ability distribution of the British workforce, in a given period, for three broadly defined occupational groups. As such, the population of interest for this research is the whole labour force of the UK. To make **inferences** about this population, we need to conduct our estimates over a random sample from that population. The BHPS was designed to be representative of the UK population at any point in time. In consequence, any derived subgroup (in this case, the labour force) should also be a representative sample of the corresponding population subgroup, in any year. However, as any longitudinal survey, the BHPS is likely to suffer from selection problems. Selection occurs when observations are missing in a non "fully random" fashion (described below), with regards to the phenomenon under study (in this case, wage determination). There are several types of selection. Here I focus on the most important ones.

1. **Missing control variables**: some workers lack data on one or several control variables needed for the estimation. Whenever these missing values cannot be imputed, these observations are necessarily be dropped from the sample. If these have a non-random pattern with respect to the overall sample (e.g. self-employed lacking more data than employees), there is selection.

2. **Labour market self-selection**: by necessity the estimation of a Mincer equation considers only those individuals who are working. However, to work is – to a certain extent – a choice, which might also depend on individuals' characteristics. If this is the case, those observed in the labour market might not be a representative sub-group of the labour force.

3. **Attrition**: this arises when survey participants leave the sample, either due to involuntary factors (like moving home, or death) or voluntary (like refusal to be interviewed). Many of those leaving are still part of the relevant population (e.g. they are working), and their exclusion might lead to an unrepresentative sample.

Selection issues do not necessarily invalidate the consistency (or internal validity) of the estimation results. Yet they could render inferences to the wider population of interest invalid. It is crucial then to understand when selection is a problem, and how to adjust for it – if possible. After this, we can move on to study the presence of selection in the sample under study.

G.1 Types of missing data scenarios

Following Wooldridge (2010), define the **population** dataset of interest as $\{y, z, u\}$, where y is the dependent variable, z are observable independent variables (endogenous or exogenous), and u are unobservable variables (to the econometrician). We observe a sample of N individuals, y_i, z_i, $i = 1, \ldots, N$. Moreover, there exists a selection mechanism s_i such that $s_i = 1$ if individual is observed in the sample, and $s_i = 0$ if not.

The most simple case of missing data is that of **missing completely at random**, or MCAR, under Rubin (1976) terminology. This occur when selection is neither a function of observable nor unobservable variables, but it happens in an unsystematic, random fashion. MCAR is defined as:

$$P(s = 1|z, u) = P(s = 1)$$

The selection mechanism in this case is: $s = f(\nu)$, where ν is a random process. Under MCAR, selection is not an issue. Estimations are consistent and valid for inferences.[1] This is of course a strong condition. It is reasonable to expect that workers with low education are more likely to be unemployed, or that those who work long hours are more likely to refuse participation in surveys. A weaker condition is known as **missing at random**, or MAR. This occur when selection is a function of observables only, as in the two examples given above (selection depending on education and working hours). MAR is defined as:

$$P(s = 1|z, u) = P(s = 1|z)$$

[1]Naturally, due to missing observations, estimates are less precise.

The selection mechanism in this case is: $s = f(z, \nu)$. Under MAR, estimations – including RE and FE models – are still consistent (Wooldridge, 2010). However, to obtain valid population inferences, selection has to be accounted for. Since selection is on observables, this is usually not a problem. The two main methods to address MAR data issues are imputation (for the case of missing control variables), and the use of sampling weights (when the individual in not sampled altogether). Many longitudinal surveys (including BHPS) contain own estimations of these sampling weights in order to facilitate researchers to account for attrition.

The third option is when some factors that are relevant for y but unobserved (i.e. u) are also important for selection. This is the case of data **missing not at random**, or MNAR. For example, selection might depend on both observables and unobservables which are informative for the wage determination (e.g. age and motivation, respectively). Unless selection is accounted for, estimations under MNAR are inconsistent, and inference is invalid. Notwithstanding, it is crucial to emphasise here that in a FE framework, selection can depend on time-invariant unobserved components (like ability) without affecting the consistency of the estimates. This is only true if there are no time-varying unobserved components (e.g. motivation again). When data is MNAR, one solution is to model both the "structural equation" and the "selection equation", as in the popular Heckman selection model, proposed by Heckman (1979). Notice that using imputation or sampling weights does not restore consistency of estimations under MNAR.

Given this framework, we can think of each selection problem separately.

G.2 Missing control variables

There are simple tests available to evaluate whether missing controls are a case of MCAR or MAR. Following Allison (2002), a t-test comparing the equality of means for the dependent variable between those with and without missing controls is rejected at 1%. As it turns out, the former have a significantly lower average wage compared with the latter. Another test – akin to an instrument exogeneity test in the RCT literature – is to run a logistic regression for an indicator variable stating whether controls are missing or not, on the regressors of interest. This test indicates whether any of these regressors are correlated with the probability of having missing controls. Evidence indicates that for the sample period under study, women and workers with less experience or lower education tend to have more missing data. Other variables like region and industry also correlated with missing controls. Together these tests indicate missing variables are not MCAR.

Imputation methods are a common solution for achieving external validity (i.e. valid inferences) when data is MAR. However, since they are conducted using information on other (observable) regressors, they are not valid under MNAR. Lacking a formal method to test whether missing controls are MNAR, in this study I use simple methods to impute data when it is "logical to do so". This is the least invasive method, and the least prone to errors, albeit time consuming. It simply consists of using information on other variables in order to impute almost sure values to those missing. For example, if region is missing in a given year, and the person reports not moving house during the last year, the previous measurement can be rolled forward. Or, filling a missing labour market status when the person reports other labour market indicators like being in a full-time job, or earning a wage. The extent of imputation performed was relatively little – never higher than 1% of non-missing values, for any regressor. After this adjustment, there is still a mild level of missing data. Considering the final sample size used in this study (90,780 observations), there are around 6,700 observations (7.4%) lacking only the union variable, 4,300 (4.8%) lacking only race, and around 3,400 (3.5%) lacking only industry or firm size. The rest of missing are rather insignificant.

G.3 Labour market self-selection

Labour market selection – particularly among women – has been a standard concern in Mincer equations since the seminal work by Heckman (1976, 1979).[2] A quick look into self-reported labour market status in the sample reveal this might be an issue here too, as Table G.1 reports. As a proportion of those who are not working, women are found mainly performing family care (45%), with full-time further education coming

[2]For an early survey of the literature, see Killingsworth and Heckman (1987). Examples of recent works include Dustmann and Rochina-Barrachina (2007) and Jacobsen, Khamis and Yuksel (2015).

Table G.1: Labour market status of sample, period 1991-2008

Labour market status	**Male**		**Female**	
	Frequency	Percent	Frequency	Percent
Unemployed	5,578	32.20	3,272	10.75
Maternity leave	20	0.12	1,730	5.68
Family care	656	3.79	13,638	44.81
Full time study	5,890	34.00	7,053	23.17
Long term sick or disabled	4,474	25.83	3,997	13.13
Govt. training scheme	325	1.88	245	0.81
Other	381	2.20	499	1.64
Total not working	17,324	100	30,434	100
Total working	68,781		63,427	

second (23%). Unemployment only comes fourth (11%). Conversely, very few male are classified under family care. Education, unemployment, and sickness are the main status for out of labour market male. Also, the number of women out of work almost as twice the number of men, although they are fairly similar when looking at those employed. If we assume that unemployment is less than a choice compared with family care, these statistics tells us that women might be more prone to selection issues. Yet, recall that selection under MCAR or MAR are not necessarily a reason of concern. This is the case if, for example, women caring for family members are a random group of the population (i.e MCAR), or they can be predicted based on recorded variables only like age, or education level (i.e MAR). Or, as already mentioned, if women's selection depend on time-invariant unobserved components (like ability) and the estimation of the wage equation is carried out using FE. In any case, a proper evaluation of self-selection requires a formal model and test, as carried out next.

Heckman selection tests

A Heckman selection model is applied to women in the sample. Define s_{It} as a selection indicator, equal to one if the individual at cluster $I(i, j)$ is in employment, and zero otherwise. The Heckman model is based on two equations. First, the "structural equation", which in our case is the wage model in equation (3.3). This is defined only for for those women in employment (i.e. $s_{It} = 1$). Second, the "selection equation", which describes the mechanism under which female self-select into employment. This mechanism is:

$$s_{It} = 1[W_{It}\gamma_t + \bar{W}_I\delta_t + \nu_{It} > 0], \quad t = 1, ..., T \tag{G.1}$$

where $1[\cdot]$ is an indicator function equal to one if the expression inside is larger than zero, and zero otherwise. W_{It} is the set of regressors that underlays women's participation decisions, $\bar{W}_I$ is the cluster mean of these regressor, and $\nu_{It} \sim$ Normal(0,1).[3] W_{It} can include variables used in the "structural equation", but it **must** include other variables not directly relevant for wages.

Here I implement two tests. The first one is a standard selection test, where the structural equation is equation (3.3). The "selection equation", outlined above, uses age, race, marital status, region, a dummy indicating whether her partner has a job, the number of hours and pay of her partner's job, and the number of children at ages 0-4, 5-11, and 12-18. The test, based on Wooldridge (2010), goes as follows:

1. Estimate equation (G.1) using Pooled Probit. This pooling assumes $\gamma_t = \gamma$ and $\delta_t = \delta$.

2. From above, calculate the inverse Mills ratio $\hat{\lambda}_{It}$ of the fitted values.

3. Using FE, estimate equation (3.3), adding the predicted $\hat{\lambda}_{It}$ into the regression.

4. Perform a Wald test to the null that $\hat{\lambda}_{It}$ is insignificant. Under the null there is no selection.

5. If selection is not rejected, obtain robust standard errors use bootstrapping methods.

The second test allows for education to be endogenous. It is based on Semykina and Wooldridge (2010), and goes as follows:

[3] The original model does not include the cluster means. However, the references from there the tests are taken (see below) includes them, following (Mundlak, 1978).

1. Estimate equation (G.1) using Probit, for each period t.

2. From above, calculate the inverse Mills ratio $\hat{\lambda}_{It}$ of the fitted values, for each period t.

3. Using FE-2SLS, estimate equation (3.3), adding each periods' $\hat{\lambda}_{It}$ into the regression. Thus, there are T number of added terms.

4. Perform a Wald test to the null that $\hat{\lambda}_{It}$ are **jointly** insignificant. Under the null there is no selection.

5. If selection is not rejected, obtain robust standard errors use bootstrapping methods.

The results of these estimations are presented in Table G.2. First, it is always informative to contrast the results of pure RE and FE models. This is shown in the first two columns. A standard Hausman test strongly supports the existence of unobserved heterogeneity. This is positive, since the purpose of this work is to estimate that heterogeneity (likely, ability). Moreover, the positive bias of education in RE indicates unobserved heterogeneity might be reflecting ability, positively correlated with education. Second, estimates accounting for selection show little evidence in favour of it. Both models using a Heckman selection mechanism reject the selection test at 10%. Third, tests from first and second stage of FE-2SLS estimation indicate the instruments (partner and parent's education) are highly valid and relevant (both at the 1%), without evidence of excluded instruments. Nonetheless, there is no evidence of endogeneity of education. These results combined make the FE model the best one.

In conclusion, while in principle it is reasonable to expect selection (specially considering the descriptive statistics in Table G.1), evidence here says the contrary. This is not the first study that finds no evidence of selection. Bonjour et al. (2003) produce a careful study of the UK labour market using data on twins, also finding no evidence of labour market selection for females. This research is of particular interest because twins data allow to directly control of ability and family background, reducing endogeneity problems.

In consequence, in what follows I assume that self-selection in the female population is not an issue. I light of this result, also I assume self-selection in the male population is even much less of a concern, and proceed to estimate all regressions that follow using both male and female population together (albeit allowing for gender-based coefficients).

G.4 Attrition

There is already evidence in the literature that the BHPS sample faces attrition issues. For example, Uhrig (2008) estimates a survival model, where the probability of attrition in a given year depend on covariates. He focuses on both the probability of non-contact and of refusal. Using BHPS data from 1991 until 2004, he finds that the likelihood of non-contact is related to having "very high" and "very low" income, being on rental and flat-type of accommodation, being unemployed, have language barriers, and having "very poor" health. Similarly, likelihood of refusal is higher for those with "very high" and "very low" income, lower education, unemployed, and with "very poor" health. From a comparative perspective, Lipps (2009) finds that BHPS has the strongest attrition at the household level, compared with other two European panel surveys.

Recall from our previous discussion that attrition leads to serious problem under MNAR. The question that follows is: are Uhrig (2008) findings an indication of MNAR? This is only the case if (i) some of the elements he identifies are indeed relevant for wage determination, and have not been modelled (for example, health); and (ii) these factors are time-varying (FE eliminates time-invariant unobservables). It is natural then to perform a test of selection due to attrition based on the MNAR assumption. Just as for the case of self-selection, the method to perform this is using a Heckman selection model.

The tests implemented here follows Wooldridge (2010). As before, two tests are conducted – assuming education is exogenous and endogenous. The mechanics of these tests mimic the ones in the previous section, and thus are not reproduced here. A minor difference is the use of first differences (instead of FE), to better capture the sequential nature of attrition. Regarding the variables in W_{it}, and following Wooldridge

Table G.2: Female labour market self-selection, 1991-2008

	Model			
Variable	**RE**	**FE**	**Heckman-FE**	**Heckman-FE-2SLS**
Experience	0.0658***	0.0779***	0.0758***	0.0772***
	(0.005)	(0.006)	(0.007)	(0.006)
Experience^2	-0.341***	-0.253***	-0.245***	-0.254***
	(0.041)	(0.058)	(0.066)	(0.035)
Education	0.0719***	0.0136***	0.0165***	0.0464
	(0.002)	(0.004)	(0.005)	(0.039)
White	0.0695**	0.0473	-0.0869	-0.545***
	(0.032)	(0.203)	(0.408)	(0.187)
Full-time	0.131***	0.124***	0.115***	0.105***
	(0.011)	(0.012)	(0.014)	(0.007)
Hours of work	0.0328***	0.0300***	0.0305***	0.0305***
	(0.001)	(0.001)	(0.001)	(0.000)
Public Sector	0.0636***	0.0348***	0.0262**	0.0258***
	(0.009)	(0.011)	(0.013)	(0.008)
Trade Union	0.0679***	0.0469***	0.0403***	0.0420***
	(0.007)	(0.009)	(0.011)	(0.006)
Married-Couple	-0.00458	-0.00591	0.0101	-0.0305
	(0.013)	(0.019)	(0.045)	(0.027)
Divorced-Widowed	-0.0130	0.0271	-0.00530	-0.0383
	(0.039)	(0.066)	(0.088)	(0.055)
Tests: (p-values in parenthesis)				
Hausman (χ^2)	355.4***			
	(0.000)			
Selection (χ^2)			0.01	17.9
			(0.933)	(0.463)
Endogeneity (χ^2)				1.04
				(0.301)
Adjusted R^2	0.654	0.489	0.495	0.433
AIC		-20518.2	-18009.4	-15867.0
BIC		-20200.1	-17690.8	-15559.2
N	31,924	31,924	26,094	24,367

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: standard errors in parentheses. All models include firm size, industry and region controls. The selection test is a Wald test on the null that the inverse Mills ratio $\hat{\lambda}_{it}$ is zero. Heckman corrected models show bootstrapped standard errors with 400 replications. The 2SLS model uses education of partner as instrument.

(2010), I use the first two lags of x_{it}, and the second and third lag of y_{it}. The latter tolerates first-order serial correlation in the idiosyncratic error of the original model. Finally, W_{it} also includes three variables which are considered important in the attrition literature: housing tenure (whether rented or owned accommodation), the number of people eligible for interview (the more to be interviewed, the more likely the possibility of refusal or selective responses), and whether the respondent is in full time work or study, part-time work or study, or neither (expected to be related to the likelihood of non-contact or refusal). Regarding the case of endogenous education, Wooldridge (2010) suggest using a subset of W_{it} as instruments. In my case, I use the first and second lag of x_{it} – excluding education – and the second and third lag of y_{it}.

Results of these exercises are presented in Table G.3.[4] Evidence of attrition is clear in both models, as then null is rejected at the 5%. There are still concerns with the 2SLS model, as the Sargan-Hansen over-identification test is rejected at 10%, meaning instruments might be not valid. They seem to be relevant though, as a test of under-identification of the instruments is rejected at 1%. A RESET test developed by Pesaran and Taylor (1999) is also implemented. The null that the model has no non-linearities omitted is rejected at 5%. All this reduces the confidence on the FD-2SLS estimation.

As said before, the test implemented above focuses on attrition as an absorbing state. This is a simplification because some workers leave the sample in a period but then return on another. As such, the extent of attrition might be much more stringent that what it is actually needed. Additionally, several years along the period the BHPS was extended by a new sub-sample in order to maintain or increase its representativeness over time. In particular, there were permanent additions in 1999 (with new data for Scotland and Wales) and in 2001 (for Northern Ireland). As such, the degree of attrition problems that existed with respect to 1991 might be different from those in the 2000s. Actually, performing the same tests described above but for the 2001-2008 period does not reject attrition at the 5% (tests are still rejected at the 10%). Yet again, this is the same strict test than before, which does not allow individuals to return to the sample.

[4]Notice that the number of observations is much limited in this model. This is because the tests implemented here can only evaluate attrition as an absorbing state, this is individuals are assume to leave the sample after a missing period is recorded – treating temporary leavers as permanent leavers. Moreover, late entrants are not included, as by construction the original sample is allegedly representative of the population.

Table G.3: Attrition estimation for first difference of log of monthly wage, 1991-2008

Variable (in differences)	**Heckman-FD**	**Heckman-FD-2SLS**
Experience	0.106***	0.0567***
	(0.010)	(0.021)
Experience^2	-0.409***	-0.193
	(0.070)	(0.128)
Education	0.0086***	0.171***
	(0.003)	(0.061)
Full-time	0.197***	0.195***
	(0.010)	(0.031)
Hours of work	0.0145***	0.0162***
	(0.001)	(0.001)
Public sector	0.0094	-0.0047
	(0.008)	(0.016)
Trade union	0.0334***	0.0305***
	(0.006)	(0.011)
Test: (p-values in parenthesis)		
Attrition (χ^2)	3.08***	26.53**
	(0.000)	(0.033)
Adjusted R^2	0.142	0.039
AIC	-8093.3	-4452.9
BIC	-7677.1	-4096.9
N	16,430	10,580

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: standard errors in parentheses. All models use marital status, firm size, industry and region controls. The selection test is a Wald test on the null that the inverse Mills ratio $\hat{\lambda}_{it}$ is zero. The 2SLS model uses first, second and/or third lag of selected exogenous variables, as suggested by Wooldridge (2010).

In fact, the female self-selection test is much more flexible, as it allows women to be in the sample after a gap in employment. As the results there indicate, selection is low.

To conclude, a strict test of attrition shows clear evidence of selection problems. Less strict tests, including those for shorter periods, show a more mild result. In any case, it is important to have these results in mind when conducting inference to the wider population.