

ESSAYS ON FIRMS AND EMPLOYEE COMPENSATION

Pawel Adrjan
St Edmund Hall
University of Oxford



Thesis submitted for the degree of
Doctor of Philosophy in Economics

Trinity Term 2018

Acknowledgements

First and foremost, I would like to thank my supervisors, Brian Bell and Steve Bond, for their guidance and support throughout the DPhil. Their insightful comments and challenging questions have helped me improve my research and have broadened my perspectives on the discipline of economics. I am fortunate to have had access to their knowledge and experience during my time in Oxford.

I am also grateful to Brian for the opportunity to co-author the final chapter of this thesis. The experience of working together on a paper was stimulating and enriching. I look forward to our continued collaboration as this research develops.

A number of people have contributed to shaping this thesis by providing comments and suggestions on specific sections, based on their reading of earlier drafts or hearing presentations at seminars and conferences. They are listed at the beginning of each chapter and I am greatly appreciative of their contributions. I would also like to thank the UK Data Service for help with accessing some of the confidential data sources used in the thesis. In addition, I am grateful to the Economic and Social Research Council and to the Department of Economics for financial support.

Last but not least, I would like to thank Javier Muñoz-Basols, whose constant support and encouragement have been vital to my development as a person and as an economist. This thesis would not have been possible without you.

Abstract

This DPhil thesis is a collection of three empirical papers that study the role of firms in the UK labour market. Each chapter focuses on firms at different points in their lifecycle.

Young firms are an engine of job creation but little is known about the quality of the jobs that they offer. In Chapter 1, I use a matched employer-employee dataset to study how starting wages and lifecycle earnings of employees differ between young and mature firms. I find that young firms pay a small premium to new hires, but subsequent wage growth is better at mature firms, both within continuing job matches and when individuals change jobs. Crucially, highly-paid and stable jobs at young firms have become increasingly rare over time, as young firms themselves have become less likely to survive and attain high productivity levels—both in absolute terms and relative to mature firms over the same period. Policies that aim to stimulate job growth by encouraging the formation of new firms should therefore pay close attention to the types of firms that form.

Chapter 2 asks what determines the proportion of a firm's income that workers receive as compensation. I use longitudinal firm data from a period of substantial labour share variation to understand the firm-level determinants of the labor share of income—a question that has typically only been addressed with country- and sector-level data. Estimating a dynamic model using GMM, I find that firms with greater market power and a higher ratio of capital to labour allocate a smaller proportion of their value added to workers. Testing the impact of tangible and intangible capital on low- and high-wage firms leads to conclusions consistent with the hypothesis of capital-skill complementarity. Overall, the results suggest that firm-level drivers play a key role in the evolution of the aggregate labour share, which has declined significantly since the 1970s.

Chapter 3 co-authored with Brian Bell, focuses on mature firms and asks how wages at such firms respond to idiosyncratic firm-level cost shocks. We create a unique dataset that links longitudinal data on workers' compensation to the unexpected costs related to firms' legacy defined benefit pension plans. We show that firms are able to share the burden of such costs when a significant share of their workers are current or former members of the plan. We also find that firms that respond to deficits by closing down the pension plans effectively reduce the total compensation of plan members. These results point to significant frictions in the labour market, which we show are a direct result of the pension arrangement that workers have. Yet closing schemes has an implicit cost for firms, since it reduces the frictions that workers face, and increases mobility.

Table of Contents

Acknowledgements	iii
Abstract	v
List of Figures	ix
List of Tables.....	xi
Introduction	1
Chapters	
1. Risky business? Earnings prospects of employees at young firms	5
1.1. Introduction.....	7
1.2. What is a young firm?	14
1.2.1. Employee data	15
1.2.2. Employer data	15
1.2.3. Matching employers with employees.....	16
1.2.4. Defining firm age	16
1.2.5. Characteristics of young and mature firms and their employees	18
1.3. Comparing wages at young and mature firms	19
1.3.1. Distribution of wages at young and mature firms	20
1.3.2. Explaining the wage differences between young and mature firms ...	21
1.3.3. Sorting and selection on unobservables	24
1.3.4. Earnings prospects of employees at young firms	28
1.4. Which young firms create high-paying jobs?	32
1.4.1. Firm success and employee wages	32
1.4.2. Where in the economy are the high-paying jobs at young firms?	36
1.5. Declining quality of young firms and real wage stagnation	40
1.6. Conclusions.....	46
1.7. Bibliography.....	48
1.8. Figures and tables.....	52
1.9. Appendices	73
2. The mightier, the stingier: Firms' market power, capital intensity, and the labour share of income.....	111
2.1. Introduction.....	113
2.2. Measuring the labour share of income at the firm level	120
2.3. Why might capital intensity and market power matter for the labour share.	124
2.4. Data	126
2.5. Empirical strategy	130

2.6. Results: Firm-level determinants of the labour share of income	134
2.7. Human capital, labour market power, and the labour share.....	142
2.8. Conclusions	149
2.9. Bibliography.....	151
2.10. Figures and tables.....	154
2.11. Appendices	163
3. Pension shocks and wages	185
3.1. Introduction.....	187
3.2. Data and empirical strategy.....	192
3.3. Firm-level cost shocks and wages.....	197
3.4. Impact of scheme closure.....	200
3.5. Frictions and worker mobility	201
3.6. Conclusions	204
3.7. Bibliography.....	206
3.8. Figures and tables.....	208
3.9. Appendices	220
Conclusion.....	227

List of Figures

1.1	Average wage gap between young and mature firms.....	52
1.2	Distribution of hourly wages by firm age	53
1.3	New hire wage premium at young vs. mature firms	54
1.4	New hire wage premium at young vs. mature firms (selection on unobservables).....	54
1.5	Returns to joining a young firm	55
1.6	Returns to joining a successful young firm	56
1.7	Where are the “good” jobs at young firms?	57
1.8	Wages by firm age over time	58
1.9	Young firm wage premium/penalty by year.....	59
1.10	Shifts in the sectoral composition of young firms.....	60
1.11	Shifts in the occupational composition of jobs at young firms	61
1.12	Rising exit rates of young firms	62
1.13	Falling productivity of young firms	62
1.14	Young firm wage premium/penalty by year and ex post firm success	63
1.A1	Firm count in the dataset vs. published aggregate time series	76
1.A2	Entry rate in the dataset vs. published aggregate time series	77
1.A3	Sectoral composition of firms in the BSD and in the matched sample	77
1.A4	Average earnings in the matched sample vs. published aggregate time series ..	78
1.B1	Contribution of young and mature firms to job creation and job destruction ...	80
1.B2	Net job creation by young and mature firms through the cycle	85
1.B3	Proportion of jobs destroyed by firm exits	88
1.B4	Unweighted firm exit rates	90
1.C1	Proportion of jobs with pension benefits.....	95
1.C2	Returns to joining a young firm based on weekly earnings	98
1.C3	Robustness of estimated returns to working at a young firm (continuing matches).....	99
1.C4	Robustness of the estimated young firm wage premium/penalty by year.....	101
1.D1	Distribution of hourly wages and log hourly wages by firm age	102
1.D2	Distribution of hourly wages by firm age in 2015	103
1.D3	Distribution of weekly earnings and hours by firm age	104
1.D4	Correlation between wages and firm size.....	105
1.D5	Age distribution of individuals with and without a prior observed wage	106
2.1	Labour share trends since 1970	154
2.2	Labour share trends over the sample period, 2005-2012	154

2.3	Labour share vs. capital intensity and market share.....	155
2.C1	Alternative measures of capital intensity and market share	174
3.1	Active members of defined benefit pension schemes, 1979-2015.....	208
3.2	Aggregate normal and special pension contributions, 1992-2015	208
3.3	Components of total compensation of employees in the UK, 1972-2015.....	209
3.4	Economy-wide vs. sample DB contributions, 2003-2015.....	210
3.5	Deficit contributions: our definition vs. company reports, 2002-2015	210
3.6	Distribution of deficit payments per worker in the sample	211
3.7	Scheme closure and employer contributions	212
3.8	Scheme closure and wages	212

List of Tables

1.1	Summary statistics: young firms	64
1.2	Summary statistics: mature firms	65
1.3	Comparing hourly wages at young and mature firms	66
1.4	Four approaches to addressing potential selection	67
1.5	Do young firms compensate employees for high exit risk? New hire wage premium.....	67
1.6	Do young firms compensate employees for high exit risk? Wage returns.....	68
1.7	Wages and ex post success of young firms	69
1.8	Start-up job creation, wages, and firm outcomes by sector.....	70
1.9	Contribution of young firms to the aggregate real wage decline	71
1.A1	Sample construction	74
1.A2	Analysis of non-response in ASHE	79
1.B1	‘Excess’ contribution of young and mature firms to job creation.....	82
1.B2	‘Excess’ contribution of young and mature firms to net job creation	86
1.B3	Job instability at young firms	91
1.C1	Robustness to alternative measures of employee compensation.....	93
1.C2	Robustness to alternative definitions of the sample	94
1.C3	Accounting for firm fixed effects	97
1.D1	Moments of the wage distribution.....	107
1.D2	Alternative methods of controlling for firm size.....	108
1.D3	Productivity distribution by firm age	108
1.D4	Wages of new hires as a predictor of a young firm’s success	109
2.1	Labour share dispersion.....	155
2.2	Summary statistics.....	156
2.3	Static models of the labour share	157
2.4	Dynamic models of the labour share	158
2.5	Robustness to alternative empirical specifications.....	159
2.6	Alternative measures of the labour share, capital intensity, and product market power.....	160
2.7	Understanding how capital intensity impacts the labour share	161
2.8	Human capital, labour market power, and the labour share	162
2.A1	Autoregressive models of the main variables of interest	165
2.A2	Reduced-form equations for 2012.....	166
2.B1	First-stage regressions	171
2.B2	Tests of overidentifying restrictions.....	172

2.B3	Additional tests of overidentification	173
2.C1	Testing the validity of using the levels equation in System GMM	175
2.C2	High- and low-wage firms.....	176
3.1	Summary statistics.....	213
3.2	Pension deficits and workers' wages.....	214
3.3	Pension deficits and DB membership	215
3.4	Scheme closure and employer pension contributions	216
3.5	Scheme closure and wages	217
3.6	Pension scheme membership and worker mobility	218
3.7	Pension scheme closure and worker mobility	219
3.A1	Sample observations by year	221
3.A2	Description of key firm-level variables	221
3.A3	Description of key worker-level variables	222
3.B1	Testing the effect of the 3 rd lag of deficit payments on wages.....	223
3.B2	Pension deficits and DB membership using lagged employment	224
3.B3	Scheme closure and employee compensation	225

Introduction

Wage and employment decisions made by firms – whether in the course of ongoing operations or in response to shocks – affect a large segment of the economically active population. In the UK for instance, employees of private-sector firms represented 64% of the labour force in 2016. A comparable figure for the US was 76%. These employment patterns beg important questions about the role of firms in the labour market. How do firms adjust wages and hiring in reaction to unexpected cost increases? How does a firm’s exposure to outside market conditions affect the share of overall value added that workers are able to capture? And what impact does the creation of new firms and firm exit have on the earnings prospects of employees? My research contributes to addressing these questions by assessing the impact of firm characteristics and firm-level events on wages and employment, using microdata on firms and individuals from the UK.

This thesis consists of three chapters that consider firms at different points in their lifecycle. Chapter 1, entitled **“Risky Business? Earnings Prospects of Employees at Young Firms,”** is motivated by the observation that young firms are an engine of job creation in many developed economies, but little is known about the quality of the jobs that such firms create. I use a matched employer-employee dataset to study how starting wages and lifecycle earnings of employees differ between young and mature firms in Great Britain. I find that young firms pay a small premium to new hires, but subsequent wage growth is better at mature firms, both within continuing job matches and when individuals change jobs. These results are confirmed by several approaches to addressing sorting and selection of employees into firms of different ages. Moreover, there is substantial

heterogeneity of outcomes: the few young firms that survive and become highly productive pay higher wages to employees from the outset than less successful young firms.

Crucially, the characteristics of young firms and the nature of the jobs that they offer have deteriorated since the mid-2000s. I document a worsening of young firms' survival rates and productivity outcomes, both in absolute terms and relative to mature firms over the same period. These trends coincide with a significant decline in relative wages at young firms. Highly-paid and stable jobs at young firms have thus become increasingly rare over time. The findings of the first chapter thus generate a nuanced assessment of the contributions of young firms to job creation, which have attracted the attention of policymakers in many countries. The implication is that policies that aim to stimulate job growth by encouraging the formation of new firms should pay close attention to the types of firms that form as a result.

As firms age and mature, they need to continually re-optimize employment and wage levels in response to competitive pressures and shocks. The next two chapters investigate how such phenomena can impact wages.

In Chapter 2, **“The Mightier, the Stingier: Firms’ Market Share, Capital Intensity, and the Labour Share of Income,”** I study the firm-level determinants of the share of employee compensation in value added (the “labour share”). The labour share began to decline in many major economies in the 1970s after decades of relative stability, generating substantial interest in the economics literature and among policymakers. Using data from a large panel of UK businesses, I document wide dispersion of labour shares across firms within narrowly-defined industries and discusses the measurement issues associated with using firm data to examine factor shares of income. I show that the dispersion in labour shares is driven by within-industry, between-firm differentials in both

average wages and productivity. To find the drivers of this dispersion while taking account of potential endogeneity problems, I estimate a dynamic model using the Generalised Method of Moments. I find that firms that attain greater market power and employ a higher ratio of capital to labour dedicate a smaller proportion of their income to compensating workers. These results highlight the importance of firm-level drivers in the evolution of aggregate income shares.

Mature, established firms in the UK (and the US) have increasingly had to face large, unanticipated increases to the costs of historic defined-benefit pension plans. Since the turn of the century, substantial deficits have arisen in many of those plans, due to extended periods of low interest rates and low market returns, as well as the increasing longevity of pensioners. UK firms, for instance, have spent £167bn on additional pension contributions since 2000, over and beyond the regular, anticipated funding costs.

Chapter 3, “**Pension Shocks and Wages**,” co-authored with Brian Bell, examines how wages at the largest firms in the UK respond to unanticipated firm-level cost shocks related to historic defined-benefit pension plans. For this purpose, we create a unique dataset that links the compensation of individual employees to the additional payments that firms have been forced to make to address their pension deficits. We show that firms are able to share the burden of such costs when a significant share of their workers are current or former members (and hence future beneficiaries) of the plan. We also investigate how compensation responds to the closure of defined benefit plans to future benefit accrual. We find that firms are able to use such closures to effectively reduce total compensation of workers who are plan members. These results point to significant frictions in the labour market, which we show are a direct result of the pension arrangement that workers have. On the other hand, closing down the defined-benefit pension schemes increases worker mobility. This action, therefore, has an implicit cost for firms since it reduces the frictions

that workers face. The chapter thus adds to a growing literature on the importance of firm-specific shocks to wages and highlights, yet again, the shortcomings of the standard competitive model for the labour market. It shows that frictions are important in understanding how workers are compensated and how firms deal with shocks.

The remainder of this document is organised into three chapters and a conclusion. Each chapter is followed by a bibliography, figures, tables, and appendices. The conclusion outlines several directions for future research that arise from this thesis.

Chapter 1

Risky Business?

Earnings Prospects of Employees at Young Firms*

Abstract

Young firms are an engine of job creation, but little is known about the quality of the jobs that they offer. I use a matched employer-employee dataset to study how starting wages and lifecycle earnings of employees differ between young and mature firms. I find that young firms pay a small premium to new hires, but subsequent wage growth is better at mature firms, both within continuing job matches and when individuals change jobs. These results are confirmed by several approaches to addressing sorting and selection of employees into firms of different ages. There is substantial heterogeneity of outcomes: the few young firms that survive and become highly productive pay higher wages to employees from the outset than less successful young firms. Yet, crucially, the characteristics of young firms and the nature of the jobs that they offer have deteriorated since the mid-2000s. I document a worsening of young firms' survival rates and productivity outcomes, both in absolute terms and relative to mature firms over the same period. These trends coincide with a significant decline in relative wages at young firms. Highly-paid and stable jobs at young firms have thus become increasingly rare. Policies that aim to stimulate job growth by encouraging the formation of new firms should therefore pay close attention to the types of firms that form as a result.

Keywords: Start-ups; Young Firms; Wages; Entrepreneurship

JEL Codes: J21, J23, J31, L26

*I thank my advisors, Brian Bell and Steve Bond, as well as Udo Brixy, Matthias Heinz, Thomas Hellmann, Ines Lee, and participants at the International Workshop on Establishment Panel Analyses (Nürnberg), 12th Conference of the Spanish Association of Labour Economists (Valladolid), Applied Micro Workshop (Oxford), and Gorman Student Research Workshop (Oxford) for helpful comments and suggestions. This paper contains data from the UK Office for National Statistics (ONS) supplied by the UK Data Service (UKDS), which are Crown copyright and are reproduced with the permission of the Controller of Her Majesty's Stationery Office and the Queen's Printer for Scotland. This work uses research datasets that may not exactly reproduce National Statistics aggregates. ONS and UKDS bear no responsibility in relation to the interpretation or further analysis of the data. All results have been reviewed to ensure that no confidential information is disclosed.

1.1 Introduction

Young firms contribute significantly to aggregate job creation. Between 2002 and 2015, firms aged three years or less accounted for an average 26% of annual job creation in the United States, calculated at the level of establishments (US Census Bureau 2015). This figure is remarkable, given that firms in this age category account for less than 10% of the total employment stock in any given year. Similar patterns have been observed in other countries and time periods (Haltiwanger, Jarmin, and Miranda 2013; Criscuolo, Gal, and Menon 2014; Anyadike-Danes, Hart, and Du 2015). That may be the reason why so many policies aimed at stimulating entrepreneurship cite job growth as a key motivation.¹

Although the quantity of jobs created by young firms is well-documented, we know relatively little about the *quality* of those jobs. Understanding how stable and well-compensated jobs at young firms are is vital for understanding the nature of those firms' contribution to the labour market. While the literature to date has mainly focused on the founders of young firms and the returns they obtain from entrepreneurial activity (e.g., Hamilton 2000; Michelacci and Silva 2007; Manso 2016), my focus in this paper is on the employees.

I investigate how starting wages and the lifecycle earnings prospects of employees at young and mature firms differ. Using a large longitudinal dataset that combines rich individual- and firm-level data for Great Britain from 2002 to 2015, I make two main contributions to the literature. First, I enhance earlier empirical work on this subject (Brown and Medoff 2003; Brixy, Kohaut, and Schnabel 2007; Ouimet and Zarutskie 2014; Burton, Dahl, and Sorenson 2017) by considering not only the differences in the initial wages of

¹For instance, fuelling “high-growth entrepreneurship, innovation and job creation” is a strategic objective of the US Small Business Administration, a government agency with over \$100 billion in loan guarantees to small businesses (SBA 2014). “Potential for business growth and job creation” was also a central requirement of the proposed ‘start-up visa’ programme for foreign entrepreneurs (International Entrepreneur Rule 2016). Similarly, the EU’s 2016 Start-up and Scale-up Initiative argues that “improving the ecosystem for start-ups and scale-ups in Europe will have a direct beneficial effect on jobs and growth in the EU.”

new hires at young and mature firms, but also subsequent earnings growth over a period of up to eight years. This is important for getting a complete picture of the returns to working for a young firm. Second, I use several approaches to investigate the role of unobserved heterogeneity of firms and individuals in estimating the relationship between firm age and wages. Understanding how this heterogeneity affects the wage differentials between young and mature firms is important for identifying which types of young firms offer well-paid jobs to employees. Third, I provide evidence of a deterioration in the survival and productivity outcomes of young firms over time, and I assess the impact of this trend on aggregate wages.

Britain is an especially interesting context for this study. Data from the OECD shows that compared to other developed countries, the British economy has consistently displayed very high rates of creation of “employer enterprises,” or businesses that employ individuals other than the owner—typically surpassed only by Israel and Korea (OECD 2016). Understanding the quality of jobs created by such businesses is particularly relevant in an economy with highly active firm dynamics, where young firms make a large contribution to job creation.

I document and explain several novel findings. While the average hourly wages and weekly earnings are significantly lower at young firms than at mature firms in the raw data (Figure 1.1), I show that virtually the entire difference is driven by the effects of tenure on the job. Employees of mature firms have substantially better opportunities to accumulate long tenures than the employees of risky, young firms, and this explains the higher average compensation levels in the former group. After controlling for observable and time-invariant unobservable attributes of firms, jobs, and individuals, I find that new hires at young firms obtain a 1-2% wage premium relative to new hires at mature firms. Despite this initial premium, however, subsequent wage growth is higher for individuals who join mature firms. This is the case both within continuing job matches and when individuals change

employers. The lifecycle returns to joining a young firm as an employee are thus worse in present-value terms than returns to joining a mature firm, for a wide range of discount rates.

Young firms are highly heterogeneous, and their observed lifespans and success rates vary greatly (Haltiwanger et al. 2013; Criscuolo et al. 2014; Decker, Haltiwanger, Jarmin, and Miranda 2014, 2016; Anyadike-Danes et al. 2015). This heterogeneity impacts the quality of the jobs that these firms offer to employees. I find a strong link between current wages and subsequent performance outcomes of young firms. The young firms that survive and achieve high productivity *ex post*, conditional on survival, pay higher wages to similar new hires *ex ante*, relative to young firms that do not survive or survive but attain low productivity levels.

I investigate this heterogeneity further to determine which sectors have the highest concentration of high-paying jobs at young firms. I find that the sectors where young firms are less likely to fail, more likely to attain high productivity, and more likely to pay high wages are specialised, high-skilled services, such as software and IT, real estate, architecture and engineering, and finance. However young firms in those sectors create a small overall proportion of new jobs. In fact, many jobs at young firms are found in low-wage, high-risk, and low-productivity sectors like bars, restaurants, food retail, and cleaning. Thus, while some young firms create well-paid and stable jobs, many do not. This suggests that policymakers who aim to stimulate job growth by encouraging the formation of new firms should pay close attention to the kinds of businesses that form as a result.

These conclusions are particularly relevant, because the composition and success rates of young firms have changed over time. I show that since the mid-2000s, young firms have become more concentrated in low-wage sectors and occupations, and their survival rates and productivity have declined—both in absolute terms and relative to mature firms over

the same period. These trends have coincided with a substantial deterioration in relative wages at young firms. At the beginning of my sample, in 2002, the observed wage of the average new hire at a young firm was 12% higher than at a mature firm. By 2015, it was 3% lower. I assess the relevance of these findings in the context of the aggregate trend in real wages, taking advantage of the fact that the sample reflects key features of the overall economy. Specifically, the median real hourly wage in the UK grew by 4% between 2002 and 2008 but declined by 10% between 2008 and 2015 (ONS 2017). I show that more than half of this decline can be attributed mechanically to the widening of the gap between wages at young and mature firms. These results highlight the importance of young firms, and of the jobs that they create, to the overall labour market.

From a theoretical standpoint, there are many reasons why young firms may compensate employees differently than mature, established firms. First, the short track records of young firms mean that they generally cannot rely on their reputation to attract job applicants, unless potential workers can observe the earlier track record of the entrepreneur. Second, young firms are likely to be learning about market conditions and their productivity (like in the model of Jovanovic 1982), meaning that they may be less productive, on average, than more mature competitors. Third, young firms may face greater financial constraints than established firms (Evans and Jovanovic 1989). Fourth, young firms may provide non-pecuniary benefits that mature firms do not, such as a more flexible schedule or a less bureaucratic work environment (Sørensen 2007; Koch, Pastuh, and Späth 2013). All else being equal, those factors are likely to put downward pressure on wages.

Wages at young firms may, however, also be buoyed by several offsetting considerations. For instance, young firms have higher failure rates than mature firms (Haltiwanger et al. 2013; Anyadike-Danes et al. 2015), and accepting employment there is a risky proposition.

Workers may, therefore, demand higher wages as a compensating differential to offset the risk of unemployment, particularly in markets where they cannot quickly find new jobs with similar pay. Moreover, a high risk of failure and lack of past reputation means that implicit contracts and promises of future wage growth are not credible. This implies that wages should be higher at the outset, and the wage-tenure profile should be flatter. In all, therefore, it is not clear from a theoretical perspective whether young firms should pay higher or lower wages to equivalent employees than mature firms, either upon hiring or later in the employment relationship.

My results suggest that many of these factors are at play simultaneously. I show that there is some evidence of negative selection of employees into young firms, which are, on average, less productive than mature firms. At the same time, relative starting wages at young firms are higher in riskier sectors, consistent with the theory of compensating differentials, while the wage growth profile is flat, consistent with the lack of credibility of implicit contracts at risky employers. Yet, even controlling for these characteristics and for time-invariant unobserved heterogeneity between employees, the present value of wages at a young firm is still lower than at a mature firm. These patterns suggest that other factors, such as non-pecuniary compensation or a lack of investment in general or firm-specific human capital, may also play a role in wage-setting at young firms. This points to interesting directions for future research.

Obtaining insight into the earnings prospects of employees at young firms requires dealing with three main conceptual and methodological challenges. The first relates to identifying young firms in the data. Firm entry into business registers is often a result of corporate activity, such as mergers or spin-offs, which creates new entities that are not necessarily new business ventures. The main focus of this paper is on firms that originate as

start-ups. I therefore use information on the corporate structure of firms to define firm age in a manner consistent with the US Business Dynamics Statistics database for the first time for the universe of British firms. This allows me to identify truly young firms separately from new registrations that arise from restructurings or mergers.

A second methodological challenge relates to the confounding effect of firm size on wages. Firm size matters for estimating the wage gap between young and old firms, because most new firms start out small, and small firms tend to pay observationally equivalent workers less than large firms do (Brown and Medoff 1989; Abowd, Kramarz, and Margolis 1999). I therefore investigate thoroughly the extent to which firm size helps explain the differences between the wages of employees at young and mature firms.

Third, the ability to estimate the effect of firm age on wages is hindered by potential sorting and selection of individuals into young and mature firms based on unobservable characteristics. However, the longitudinal nature of the dataset makes it possible to track employees over time as they move between young and mature firms. This enables me to explore four approaches to addressing selection issues, over and above controlling for observable characteristics of individuals, jobs, and firms. The first two approaches are based on comparing wages at young and mature firms within two specific sub-samples of employees: displaced workers and labour market entrants. This is motivated by the assumption that, in both cases, the decision of those individuals to search for a new job is not correlated with unobserved individual characteristics, unlike the voluntary decision of an employed worker to switch to a new job (see, e.g., Neal 1995; Couch and Placzek 2010; Huttunen, Møen, and Salvanes 2011). The other two approaches consist of comparing the starting wages of the same individuals as they join firms of different ages at different moments in their career, using fixed-effects estimation to control for unobservable, time-

invariant individual and firm characteristics. I do this both for the sub-sample of labour market entrants, who can be tracked over time, and for the entire sample of employees. While each of these approaches has its advantages and disadvantages, they all lead to highly consistent conclusions. Together, therefore, they enable me to characterise the relationship between firm age and wages in ways that have not been addressed in the existing literature.

My results contribute to an emerging body of empirical work on the compensation of employees at young firms. Existing research has reached mixed conclusions about whether starting wages at young firms are higher or lower than at mature firms. Brown and Medoff (2003) found that the positive relationship between firm age and wages disappeared after controlling for worker characteristics in a cross-section of approximately 1,000 employees in the US. Other work, however, has found persistent wage gaps between young and mature firms. Aggregating at the firm level, Brixy et al. (2007) found an average 8% wage penalty at new German firms created between 1997 and 2001 relative to incumbent firms, after taking into account a range of firm and workforce characteristics. Wage penalties at young firms in the 1-3% range were also present in several studies that combined firm-level data with information on individuals to study specific populations, such as Swedish labour market entrants and job switchers (Nyström and Elvung 2014, 2015). On the other hand, Heyman (2007) found that the conclusions about relative wages at young and mature firms differed depending on the period studied. Furthermore, Ouimet and Zarutskie's (2014) firm-level analysis for the US from 1992 to 2004 concluded that young firms paid a premium to new hires, compared to mature firms. Similarly, Burton et al. (2017) found that young Danish firms observed between 1991 and 2006 paid a premium to new full-time hires, even after accounting flexibly for a wide range of observed attributes of firms and individuals.

One potential reason why the conclusions of these studies are contradictory is that many

of them focus on specific populations or pay limited attention to unobserved characteristics that affect the selection of workers into firms of different ages. In contrast, in this paper, I use several approaches to address potential selection on unobservables. I show that neglecting to account for potential selection leads to different conclusions about the wage differences between young and mature firms. In addition, I complement existing work by considering wage growth over time. These analyses shed a new light on the nature of job creation at young firms.

The paper is organised as follows. Section 1.2 describes the data sources and the construction of the firm age variable. Section 1.3 documents the differences in the distribution of wages at young and mature firms. It investigates the contribution of firm, job, and individual characteristics to explaining the observed differences in wages between young and mature firms, and assesses the lifecycle earnings prospects of the employees who join those firms. Section 1.4 assesses how starting wages at young firms vary depending on the firm's ex post success. This section also evaluates in which segments of the economy high-paying, stable jobs at young firms can be found. Section 1.5 then investigates the relationship between the deteriorating qualities of young firms over time and the relative decline in wages at young firms, compared to mature firms. Section 1.6 concludes.

1.2 What is a young firm?

To study the characteristics of jobs at young firms and the relationship between firm age and wages, I match annual data on a sample of employees to the administrative register of firms in Great Britain from 2002 to 2015. This section describes the data sources, the construction of the firm age variable, and the characteristics of the final sample used in the empirical analysis.

1.2.1 Employee data

Employee data is taken from the Annual Survey of Hours and Earnings (ASHE) administered by the Office for National Statistics (ONS). ASHE is a 1% random sample of all employees in Great Britain registered on the national Pay As You Earn (PAYE) income tax and social security system. Data are provided by employers, who are legally required to supply it in April each year. As a result of its source in payroll data, ASHE provides high-quality information and serves as one of the main sources for analysing the structure and evolution of British wages.²

ASHE supplies rich data on individual compensation (including bonuses, overtime pay, and pension contributions), as well as personal and job characteristics such as age, gender, occupation, workplace location, contract type, and hours worked. Compensation is not top-coded, so the full distribution is observed. A unique personal identifier can be used to trace each individual over time since 1997. Firm identifiers are provided from 2002 onwards, allowing for employees to be matched to employers, and thus determining who works at young and mature firms.

1.2.2 Employer data

Employer data comes from the Business Structure Database (BSD), an annual April snapshot of all firms registered for value-added tax or PAYE in Great Britain. The BSD contains information on firm employment, sales, sector, registered address, legal form, and ownership since 1997.³ The dataset covers an estimated 99% of formal economic activity (ONS 2006)

²Researchers routinely use ASHE for empirical work. This dataset has been used, for instance, to study the impact of the Great Recession on real wages (Gregg, Machin and Fernández-Salgado 2014) or the effects of increases in the minimum wage (Bell and Machin 2018).

³The BSD defines a firm as an ‘enterprise,’ i.e., the smallest organizational unit producing goods and services that has autonomy over its resource allocation decisions (ONS 2006). A large corporate group may thus include several ‘firms’ linked by common ownership. I take these linkages into account when defining firm age, as described further below.

and serves as the sampling frame for business surveys used to produce GDP estimates for national accounts.

1.2.3 Matching employees with employers

Employee data from ASHE is matched to employer data from the BSD based on a unique firm reference number provided by the ONS. I restrict the sample to employees whose reported hourly wage does not reflect a junior/trainee rate and was not affected by absence. These exclusions serve to avoid potential distortions in reported wages and are standard in published aggregates derived from ASHE. In addition, I exclude sectors dominated by public employers and focus on the ‘business economy,’ composed of: mining, energy, manufacturing, construction, trade, finance, transportation, accommodation and food, real estate, renting, and business activities (SIC 2003 sectors C through K).

The resulting sample consists of 1,329,513 job-year observations on 269,258 individuals working at 131,099 firms. Of those individuals, 17,131 are observed working at both young and mature firms at different points during the sample period. They will be relevant for the estimation of results with individual fixed effects. Appendix A describes the construction of the sample in greater detail and shows that the final dataset matches several stylised facts of the overall UK business economy.

1.2.4 Defining firm age

Comparing the wages of employees at young and mature firms requires determining the age of each firm in the dataset. One challenge is that corporate events such as mergers or restructurings often lead to the creation of new firm identifiers, and it is difficult to distinguish such cases from genuine firm births. Since the main interest of this paper is in identifying the employees of firms that start as independent, new ventures, additional steps

are needed to define firm age in a way that distinguishes between different forms of entry.

To define firm age in the data, I exploit available information on each firm's corporate structure. When a firm first appears in the BSD, I set its birth year to be the birth year of its oldest establishment. Each firm is then aged by one year every year. The establishment birth year is defined as the year when the establishment was first identified as a workplace of at least one employee in the tax records. I classify firms observed in their year of birth as 'new' if they are not owned by a pre-existing corporate group, and if they enter with employment of less than 2,500.⁴ On the other hand, if a firm is observed as belonging to a pre-existing group in the year of entry into the dataset, or if it enters with employment of 2,500 or more, then I classify it as a 'not truly young' firm. Those firms are most likely entering the register as a result of mergers, spin-offs, or restructuring activity, rather than being new, business ventures. Separating out 'not truly young' firms is important for ensuring that the conclusions about wages at young firms are not contaminated by the inclusion of established businesses.

This process provides an internally-consistent measure of each firm's age over time that is not affected by the opening and closing of establishments post-entry. It avoids incorrectly classifying new entries into the register as new firms if they result from M&A activity or restructuring involving pre-existing firms. It is also consistent with the algorithm used to derive firm age in the US Business Dynamics Statistics database, as described in Becker, Haltiwanger, Jarmin, Klimek, and Wilson (2006). This makes the resulting dataset useful for cross-country comparisons of firm entry and exit activity.

The analyses in this paper compare the wages of employees at 'young' and 'mature' firms. Throughout the paper, I define young firms as those that are aged 3 years or less,

⁴This threshold is chosen to be consistent with US Business Dynamics Statistics.

and mature firms as those aged 10 years or more. These thresholds are motivated by three observations. First, Appendix B shows that the contributions to job creation are greatest for young firms aged 3 years or less, so this firm age category is the most interesting to analyse from a policy standpoint. Second, raw data on average hourly wages and weekly earnings by firm age in Figure 1.1 shows that the gradient flattens just before the age of 10, suggesting that this is a good threshold for defining a mature firm. Third, setting the threshold for a mature firm at age 10 is consistent with how mature firms are often defined in the literature (e.g., Haltiwanger et al. 2013).

1.2.5 Characteristics of young and mature firms and their employees

Young and mature firms, and their employees, differ on several observable dimensions. Table 1.1 presents summary statistics on annual observations at firms aged 3 years or less and firms aged 4-9 years. Table 1.2 presents the same data for mature and ‘not truly young’ firms.

These summary statistics show that average hourly wages and weekly earnings at young firms are lower than at mature firms. The median real hourly wage is £9.7 at firms aged 3 years or less, £10.7 at firms aged 4 to 9, and £11.2 at firms aged 10 or more. Weekly earnings rise more steeply with firm age than wages, because hours worked are higher at older firms.

Young firms tend to have characteristics that are typically associated with low wages. For instance, the average young firm is smaller than the average mature firm. Median employment is 9 for observations at firms aged 3 years or less, 20 at firms aged 4 to 9, and 1,528 at firms aged 10 or more. Median real annual sales in those three firm age categories are £0.5 million, £1.6 million, and £162.9 million, respectively. Another noteworthy pattern is that average sales increase more steeply with firm age than average employment. This

suggests that older firms generate higher sales per worker than younger firms. Jobs at young firms are also more likely to belong to low and medium skill categories than jobs at mature firms. Finally, young firms are more likely to offer jobs in the service sector, where wages are generally low.

Tables 1.1 and 1.2 also show that firms classified as ‘not truly young’ differ from young firms in two main ways. First, the average ‘not truly young’ firm is larger, with median employment of 149 and real annual sales of £9.9 million. It is therefore significantly more sizeable than the median firm aged 3 years or less (which has employment of 9 and sales of £0.5 million) or the median firm aged 4-9 years (which has employment of 20 and sales of £1.6 million). Second, the average tenure in the job of individuals working at ‘not truly young’ firms is greater than the average tenure of individuals working at young firms aged either 3 years or less or 4 to 9 years. In fact, 27% of the employees observed at this group of firms have reported tenure that is longer than the firm’s existence in the dataset. This suggests that many of those firms were existing businesses that had undergone a transformation that generated a new entry in the register. Appendix B shows that the patterns of job creation and job destruction of ‘not truly young’ firms are similar to those of mature firms. All these attributes confirm the suspicion that those are not new ventures, justifying the decision to classify them as a distinct category.

1.3 Comparing wages at young and mature firms

This section begins by using the matched sample to document differences in the observed wage distribution at firms of various ages. It then takes advantage of the longitudinal nature of the dataset to assess how the wages of new hires, and subsequent wage growth, differ between young and mature firms for similar individuals performing similar jobs, using

several methods to address potential sorting and selection based on unobservables.

1.3.1 Distribution of wages at young and mature firms

Figure 1.2 compares the distribution of real hourly wages across all jobs at firms aged 3 or less (shaded bars) and firms aged 10 or more (transparent bars). Panel A considers all employees in the dataset, pooling across all years in the sample. The chart on the left-hand side shows the main part of the distribution, up to and exceeding the 90th percentile, while the chart on the right-hand side shows most of the upper tail of the distribution, up to and exceeding the 99th percentile. Wages at young firms are substantially more concentrated at the low end of the distribution, compared to wages at mature firms. Low wages at young firms are concentrated mainly at the minimum wage.⁵ The distribution of wages at mature firms has greater mass in the middle, but, interestingly, the two distributions look relatively similar at the upper end. Thus, while many employees at young firms earn very low wages, (and hence the average wage at a young firm is lower than at a mature firm), some are compensated very highly, even above the 90th or 99th percentile of the overall distribution. Later in this paper, I shed some light on this dispersion by investigating what kinds of young firms pay high wages.

The differences between the distributions of hourly wages at young firms and mature firms are far less pronounced when considering only wages paid to new hires in Panel B. In other words, the large differences in the distribution (and in the average) of wages at young and mature firms can largely be ascribed to differences in employees' tenure on the job.⁶ Tenure is mechanically low at a young firm, since an employee cannot have been employed at the firm any longer than the firm has been in existence. In contrast, employees at mature

⁵Figure 1.D2 in the Appendix illustrates this using the last year of the sample period.

⁶Figure 1.D3 in the Appendix shows that the same conclusions hold for weekly earnings.

firms have a wide range of tenures, from one day to several decades. At the same time, tenure is associated with higher wages. Tenure is thus probably the reason why studies based on payroll data aggregated at the firm level find that young firms pay substantially less than mature firms (e.g., Brixy et al. 2007, Ouimet and Zarutskie 2014).

1.3.2 Explaining the wage differences between young and mature firms

The fact that the wage distribution is similar at young and mature firms does not necessarily mean that a given employee would be compensated equally at firms of different ages. Young firms and mature firms may differ on a number of dimensions that could be correlated with wages. For instance, young firms might be more likely to offer jobs in well-paid occupations (such as computer programming or engineering), but due to their limited track record, they might attract relatively less productive employees than mature firms. The result could, therefore, be that the initial wage, and perhaps also the future earnings prospects, of an individual endowed with specific skills, abilities, and training depended on the age of the firm, even if the overall wage distributions were the same in the data. This section explores the extent to which employees are compensated at firms of different ages.

I focus on a simple metric: the gap in mean log hourly wages between young and mature firms.⁷ To assess first how observed characteristics of firms, jobs, and employees explain

⁷There are two reasons for focusing on log hourly wages, rather wages in levels or weekly earnings. First, the mean log wage is more representative of a typical (median) employee than the mean wage. The distribution of log wages is also closer to normal, with lower skewness that does not vary substantially with firm age. Figure 1.D1 in the Appendix illustrates this visually, while Table 1.D1 reports selected moments of the distribution of real hourly wages and log real hourly wages by firm age. Second, hourly wages represent remuneration per unit of labour input. Thus, they provide a measure of compensation that is consistent with theoretical models of wage determination. For these reasons, log wages are a standard measure of compensation in the labour economics literature. Appendix C shows that the results of this paper are robust to using weekly earnings as the dependent variable.

this gap, I estimate the empirical model of individual wages given by

$$w_{ijt} = \beta_0 + \beta_1 \mathbb{1}(\text{Firm age} \leq 3)_{jt} + \beta_2 \mathbb{1}(4 \leq \text{Firm age} \leq 9)_{jt} + \beta_3 \mathbb{1}(\text{Not truly young})_{jt} + \delta_t + X_{ijt}\Gamma + \varepsilon_{ijt} \quad (1.1)$$

which models the log real hourly wage, w_{ijt} , of individual i , observed working at firm j in year t . There are three dummy variables for firm age, with mature firms aged 10 years or more being the omitted category. The coefficients on each firm age dummy therefore represent the approximate percentage difference in wages relative to mature firms. Year dummies δ_t absorb annual trends in average wages. Finally, the vector X_{ijt} contains additional controls and fixed effects, described further below.

Figure 1.3 gives a visual summary of the results of estimating equation (1.1), pooling all observations in the sample. The height of each bar represents the estimated coefficient $\hat{\beta}_1$ on the dummy for a young firm. Each of the bars corresponds to a different specification of the vector of controls, X_{ijt} . The underlying regression results are reported in greater detail in Table 1.3.

The first bar shows the coefficient on the young firm dummy when the model is estimated on the full sample of employees, and the vector of controls X_{ijt} is empty. On average, wages at young firms are 12.6% lower (log difference -0.135) than at mature firms. However, the second bar shows that when comparing only the wages of new hires, the estimated wage penalty at young firms relative to mature firms turns into a premium of 2.5%. Tenure effects thus play an important role in explaining why wages are substantially lower at young firms than at mature firms when averaging across all jobs.

The next three bars examine the extent to which the estimated wage difference is affected by controlling for observable characteristics of firms, jobs, and employees. The third bar

continues to consider only new hires but controls additionally for firm size and sector. Firm size is a crucial control variable. Young firms typically start out small (Haltiwanger et al. 2013), and a well-established literature shows that smaller firms typically pay lower wages than large firms (Brown and Medoff 1989; Abowd et al. 1999). The limited empirical literature on wages at young firms has therefore recognised the importance of distinguishing firm age effects from size effects (see Brown and Medoff 2003 for an early discussion). I control flexibly for size by adding dummies for 19 ventiles of employment.⁸ Once controls for size and 5-digit sector are in place, the wage premium at young firms rises slightly to 3.1%.

The fourth bar controls additionally for observed attributes of the job. These controls consist of: region, 2-digit occupational code, whether the job is part-time or full-time, and whether the job contract is temporary or permanent. This leads the estimated wage premium to fall to 0.8%.

Finally, the fifth bar adds controls for observed personal attributes. The personal variables available in the dataset are age, age squared, and gender. ASHE lacks data on education or training, so I control for each individual's occupation-specific experience, estimated as the number of prior periods in which the individual was observed working in the same occupation. Occupation-specific experience is an important determinant of wages (Kambourov and Manovskii 2009, Sullivan 2010, and Zangelidis 2009 using British data) and is the best available proxy for education in this dataset.⁹ I also control for the average of the individual's past log hourly wages to capture unobservable attributes of employees that

⁸Figure 1.D4 in the Appendix explores the firm size-wage relationship in detail. Table 1.D2 shows that the results on the estimated wage gap between young and mature firms are similar, whether ventiles, centiles, or a polynomial specification are used for firm size, but they differ significantly when size is only controlled for linearly (as the literature has typically done).

⁹In Sullivan's (2010) estimates, the coefficient on occupation-specific experience is more statistically significant and larger in magnitude than that on education.

may be correlated with wages, such as individual productivity. Finally, I add a dummy for whether the individual has previously worked at a young firm to address the possibility that firms differentiate among workers based on prior experience at a firm of similar age.¹⁰

Controlling for these individual characteristics reduces the estimated wage gap between young and mature firms to a precisely-estimated zero. The results in the remainder of the paper control for all of the above attributes, which I will refer to as ‘full controls.’¹¹

1.3.3 Sorting and selection on unobservables

Controlling for observables does not address the potential bias from the sorting and selection of individuals into young and mature firms based on omitted, unobserved variables that may be correlated with wages. The direction of this bias is difficult to determine. Consider two scenarios. First, suppose that individuals who are naturally risk-averse are attracted to stable jobs at mature firms, and that risk aversion is negatively correlated with wages. This could happen, for instance, if risk-averse individuals have a lower reservation wage (Cox and Oaxaca 1989). In this scenario, the estimate of the wage gap will be biased upwards. Second, suppose instead that young firms attract individuals with lower ability and skills than mature firms, due to their short track record and greater instability. Since low ability and skills are likely to be associated with low wages, the estimate of the wage gap will be biased downwards. Both scenarios seem plausible. The estimates from the prior section are therefore not well-suited for drawing conclusions about the impact of firm age on wages.

In this section, therefore, I use four approaches to address the potential confounding effects of the unobservable attributes of employees. Each approach focuses on a different

¹⁰Left-censoring of individual labour market histories is mitigated by using data going back to 1997 to estimate occupation-specific experience and past wages. Controlling for these variables leads labour market entrants to be dropped from the sample. However, I consider this group separately below.

¹¹Related work on the starting wages of new hires has used propensity score matching as another approach to control for observable differences between the employees of young and mature firms (Nyström and Elvung 2014; Burton et al. 2017).

sub-sample of individuals, involving different sets of treatment and control groups, and has its own strengths and weaknesses. I report all four to build a comprehensive picture of the wage differences between new hires at young and mature firms.

The first approach is to consider displaced workers. These are individuals who were forced out of their previous job by the employer shutting down. I compare the wages of displaced workers who joined young firms with those who joined mature firms by estimating equation (1.1) with full controls on the sub-sample of displaced workers only. To the extent that firm exit is a random event uncorrelated with unobserved individual characteristics, focusing on this sub-sample can address concerns about selection into job search. This approach is often used in the labour economics literature (e.g., Neal 1995; Couch and Placzek 2010; Huttunen et al. 2011). However, there are two main reasons for caution. The first is that while focusing on displaced workers may address the “push” factor that leads employees to look for another job, it does not address the “pull” factor that leads them to choose a job at a young firm. Second, it is possible that displaced workers may be non-random if they were previously selected into riskier firms based on some unobservable characteristics.

A second, complementary approach focuses on labour market entrants, as individuals with no prior experience who could not have previously sorted into failing firms. Labour market entrants are an interesting group to study in their own right, given that many countries implement policies focused on youth employment. I define labour market entrants as individuals that are at most 25 years old when they are observed for the first time in the dataset. I then estimate equation (1.1) with full controls on the sub-sample of labour market entrants. Like in the case of displaced workers, however, the choice of whether to join a young or a mature firm could be driven by omitted, unobserved variables. The results may

thus suffer from similar biases.¹²

The third approach is to address the biases that arise from a labour market entrant's choice of employer by comparing the wages of the same individual during employment spells at firms of different ages. The empirical implementation consists of estimating equation (1.1) with individual fixed effects using the sub-sample of individuals who have been observed as entering the labour market for the first time at some point during the sample period. Two additional modifications are needed. First, the firm age dummies are interacted with an indicator of whether the employee is a new hire in a given period. Second, job tenure is added to the vector of controls to account for the effect of tenure on wages in other periods. In this approach, identification of the wage differential between young and mature firms is based on the assumption that unobservable individual attributes that affect both wages and the likelihood of being hired by a young firm are time-invariant, conditional on the full set of time-varying individual characteristics and labour market variables described above.

The fourth, and final, approach is to extend the estimation with individual fixed effects to all employees in the dataset. Here, the control group for any employee joining a young firm is the same individual joining a mature firm at another point in their career. The empirical implementation is the same as for labour market entrants with individual fixed effects, except that the sample consists of all employees observed more than once in the dataset.

Figure 1.4 summarises the results for the starting wages of new hires using the four approaches. Each bar reports the coefficient on the young firm dummy. Among displaced workers, the estimated wage premium for new hires at young firms compared to new hires at mature firms is 2.1%. This estimate is not statistically significant due to a small sample

¹²There are also other ways in which labour market entrants, defined in this manner, are a self-selected group. To be captured in the dataset, they must have chosen to enter wage work, as opposed to, for instance, self-employment or non-employment. However, this applies equally to all individuals in the dataset, given this paper's focus on employees.

size.¹³ The estimated premium for labour market entrants is 1.0%, significant at a 10% level. These results suggest that negative selection into job search may have played a role in the earlier finding of no wage difference between young and mature firms. Adding individual fixed effects to the regression for labour market entrants increases the estimated premium to 2.0%, significant at a 1% level. This suggests that negative selection of individuals into jobs at young firms may also be present, although the difference between the estimate with and without individual fixed effects is not statistically significant. Finally, extending the model with individual fixed effects to all employees gives an estimated wage premium of 1.4%. The results of all four approaches are directionally consistent with one another and suggest that young firms pay a 1-2% starting wage premium to similar new hires, relative to mature firms.¹⁴

The conclusion that young firms pay slightly higher wages to new hires than mature firms is robust to a range of alternative specifications. Appendix C shows that the conclusions are not sensitive to different definitions of a young firm, the potential presence of firm owners in the data, or alternative measures of compensation, such as weekly earnings or total compensation including fringe benefits.

The results are also robust to including firm fixed effects to control for time-invariant, unobserved attributes of employers, such as firm-specific wage policies not explained by observed firm characteristics. Adding firm fixed effects to the model with individual fixed effects and using the full employee sample increases the estimated hourly wage premium to 2.3%, significant at a 1% level (see Table 1.C3). Moreover, R^2 rises to 0.93 (compared to 0.20 without firm fixed effects), suggesting that unobserved firm attributes explain a substantial

¹³There are 588 individuals observed as joining a young firm immediately after their previous employer shut down. The estimated premium is insignificant also when using the full sample of new hires and interacting the firm age dummies with an indicator for whether the worker was displaced.

¹⁴The underlying regression results are reported in Columns (1) to (4) of Table 1.4.

part of the variation in individual compensation, consistent with the literature (Abowd et al. 1999). Therefore, later in this paper, I investigate to what extent the heterogeneity of young firms matters for wages.

1.3.4 Earnings prospects of employees at young firms

This section provides evidence on the wage growth profile of individuals who join young firms, compared to those who join mature firms. A good starting point is to consider each employment spell in the dataset and compare how wages subsequently evolve for those individuals who join a young firm relative to individuals who join a mature firm, regardless of whether those individuals remain with the same firm or switch jobs in future periods. I call these ‘unconditional returns.’ I use them further below as a benchmark for comparing the returns in continuing worker-firm matches. The unconditional returns are particularly relevant if working at a young firm leaves a lasting (positive or negative) imprint on an individual’s future earnings potential.

To assess the unconditional returns to joining a young or mature firm, I normalise the starting year of each observed employment spell to zero.¹⁵ I then estimate equation (1.2) below. This is an enhanced version of equation (1.1) with full controls and individual fixed effects, expanded to include dummies for each year k elapsed since the beginning of the

¹⁵I define a spell loosely as an observed employment relationship between a firm and a worker. Since the dataset is based on an annual snapshots, it does not capture information about spells that do not span the April reporting date. When individuals are observed as switching jobs between consecutive reporting dates, the starting date of the new job is known, but the exact end date of the previous job spell is not.

employment spell, $\mathbb{1}(\text{Years} = k)_{ijt}$, interacted with firm age dummies.

$$\begin{aligned}
w_{ijt} = & \beta_0 + \beta_1 \mathbb{1}(\text{Firm age} \leq 3)_{jt} + \sum_{k=0}^{13} \beta_{1,k} \left(\mathbb{1}(\text{Firm age} \leq 3)_{jt} \times \mathbb{1}(\text{Years} = k)_{ijt} \right) \\
& + \beta_2 \mathbb{1}(4 \leq \text{Firm age} \leq 9)_{jt} + \sum_{k=0}^{13} \beta_{2,k} \left(\mathbb{1}(4 \leq \text{Firm age} \leq 9)_{jt} \times \mathbb{1}(\text{Years} = k)_{ijt} \right) \\
& + \beta_3 \mathbb{1}(\text{Not truly young})_{jt} + \sum_{k=0}^{13} \beta_{3,k} \left(\mathbb{1}(\text{Not truly young})_{jt} \times \mathbb{1}(\text{Years} = k)_{ijt} \right) \\
& + \sum_{k=1}^{13} \theta_k \left(\mathbb{1}(\text{Years} = k)_{ijt} \right) + \delta_t + a_i + X_{ijt} \Gamma + \varepsilon_{ijt}
\end{aligned} \tag{1.2}$$

Panel A in Figure 1.5 plots the results. Returns are expressed relative to the starting wage of new hires at a mature firm, normalised at zero. Relative returns at a young firm are represented by a solid line, while returns at a mature firm are represented by a dashed line. As shown before, hourly wages are initially higher at the start of a spell at a young firm, relative to the start of a spell at a mature firm. However, subsequent real wage growth is negative at a young firm and positive at a mature firm. Unconditional wage returns at a young firm are therefore significantly lower than at a mature firm.¹⁶

One reason why unconditional wage growth after joining a young firm is lower relative to joining a mature firm is that spells at young firms are more likely to end in separation. Switching jobs is often associated with a decline in wages, even if the transition is voluntary (Jolivet, Postel-Vinay, and Robin 2006; Nosal and Rupert 2007). Moreover, many of the transitions out of young firms are, in fact, involuntary because many of those firms shut down within a few years of entry (see Figure 1.B3). Such transitions are typically associated with large and persistent earnings losses (Couch and Placzek 2010). Joining a young firm is thus, on average, risky business.

Do the expected returns improve when considering only firms that survive and continuing

¹⁶Figure 1.C2 in the Appendix shows similar patterns for weekly earnings.

worker-firm matches? Panel B of Figure 1.5 illustrates the answer by plotting the marginal effects of firm age at each year of tenure from equation (1.3) below. This is, once again, an enhanced version of equation (1.1) with full controls and individual fixed effects. This time, however, firm age dummies are replaced by dummies indicating the age of the firm when the individual was originally hired (abbreviated as ‘Age’ in the equation). These dummies are then interacted with indicators for each year k of tenure in the job, $\mathbb{1}(\text{Tenure} = k)_{ijt}$. This equation can therefore be used to estimate the wage returns at young and mature firms at different tenure points, relative to the baseline group of new hires at mature firms.

$$\begin{aligned}
w_{ijt} = & \beta_0 + \beta_1 \mathbb{1}(\text{Age} \leq 3)_{jt} + \sum_{k=0}^{13} \beta_{1,k} \left(\mathbb{1}(\text{Age} \leq 3)_{jt} \times \mathbb{1}(\text{Tenure} = k)_{ijt} \right) \\
& + \beta_2 \mathbb{1}(4 \leq \text{Age} \leq 9)_{jt} + \sum_{k=0}^{13} \beta_{2,k} \left(\mathbb{1}(4 \leq \text{Age} \leq 9)_{jt} \times \mathbb{1}(\text{Tenure} = k)_{ijt} \right) \\
& + \beta_3 \mathbb{1}(\text{Not truly young})_{jt} + \sum_{k=0}^{13} \beta_{3,k} \left(\mathbb{1}(\text{Not truly young})_{jt} \times \mathbb{1}(\text{Tenure} = k)_{ijt} \right) \\
& + \sum_{k=1}^{13} \theta_k \left(\mathbb{1}(\text{Tenure} = k)_{ijt} \right) + \delta_t + a_i + X_{ijt} \Gamma + \varepsilon_{ijt}
\end{aligned} \tag{1.3}$$

Panel B of Figure 1.5 plots these relative returns. Even when conditioning on continuing matches, employees hired by young firms subsequently experience lower average wage growth than those hired by mature firms. Focusing on the first five years and assuming the same number of hours worked per year, the present value of wages at a young firm only exceeds the present value of wages at a mature firm for a discount rate greater than 50%. Thus, even in this highly selected sample of surviving firms and continuing employees, joining a mature firm pays more in expectation than joining a young firm, for an individual who is able to choose between the two.¹⁷

¹⁷Figure 1.C3 in the Appendix shows that the patterns of lifecycle wages within continuing matches are robust to excluding foreign-owned firms and individuals who may be firm owners. Those were the most important of the potential concerns tested in the context of the initial wage premium at young firms above. When such firms and individuals are excluded from the sample, the returns to being hired at a young firm are still lower than the returns to being hired at a mature firm, both unconditionally and conditioning on continuing

These results are somewhat surprising if one believes that wages at young firms should offer a compensating differential for the high risk of job loss resulting from firm failure. One potential explanation that reconciles this view with the results would be that employees are simply unaware of the high exit rates of young firms. However, two pieces of evidence argue against this view. First, Table 1.5 shows the results of enhancing the individual wage regression given by equation (1.1) with the additional variable that measures the differential between the average 3-year exit rates of young and mature firms within each 5-digit sector over the sample period. This exit rate differential is also interacted with the young firm dummy to measure whether the new hire wage premium at young firms relative to mature firms is higher in sectors where there is a larger gap between the exit rates of young and mature firms. The results show that a 1 percentage-point increase in the sectoral exit risk differential leads to a 0.8% rise in wages at young firms relative to mature firms in that sector. Given that these results also control for individual fixed effects, they suggest that employees differentiate between risky and safe jobs in their wage demands.

Second, Table 1.6 reports the wage premium at young firms relative to mature firms at different points in time since joining the firm. The sample is split into two, based on whether the exit rate differential is below or above that of the median sector. In sectors where young firms are much more likely to fail than mature firms (column 1), the initial wage premium at young firms relative to mature firms is higher than in sectors where young and mature firms fail at similar rates (column 2). Yet, the subsequent wage profile is similar across the two columns. The net present value of the wage differential over the first 5 years of the match is

matches.

Other robustness checks (not included here) explore the extent to which returns to working at a young firm vary across the regions of the UK. In particular, the densely agglomerated economy of Greater London is often highlighted as having different characteristics than the rest of the UK, e.g., higher house prices and a large financial centre, which may play a role in wage determination (Gregg et al. 2014). Splitting the sample shows that young firms based in London (<15% of the sample) do not pay an initial premium while offering better wage growth compared to young firms elsewhere. However, the estimated present value of wages at young firms is still lower than at mature firms in London, in line with the results for the rest of the sample.

therefore higher in sectors where the relative risk of joining a young firm is higher, for any discount rate. This argues against the lack of knowledge on the part of the employees.

In the results above, returns to tenure in a continuing match at a young firm are better than the unconditional returns. Moreover, there is some evidence that at a sectoral level, wages at young firms are relatively higher when exit risk is higher. These patterns suggests that firm heterogeneity may play an important role in the earnings prospects of the employees who join young firms. The next section of the paper explores this further.

1.4 Which young firms create high-paying jobs?

The result that joining a young firm entails an initial wage premium compared to joining a mature firm, but subsequent wage growth is better at a mature firm, reflects *average* outcomes. An appropriate analogy would be to an individual—with their observed and unobserved attributes—who chooses to enter a job with specific observable characteristics, such as sector and occupation, but ultimately decides which firm to join by rolling a dice. It is therefore relevant to ask: does it matter what *type* of young firm that individual joins? In this section I assess how the earnings prospects of employees at young firms differ depending on the ex-post success of their employer, where firm success is measured in terms of survival and productivity outcomes. I then look at the sectoral composition of young firms to determine whether jobs created by those firms tend to be concentrated in sectors with particular characteristics, based on average wages, survival rates, and productivity.

1.4.1 Firm success and employee wages

Similar employees are compensated differently at different firms. One indication of this is the role of firm fixed effects in explaining the variation in individual wages and earnings of new hires, measured by R^2 . This is consistent with the literature, which finds that individuals

with the same time-invariant characteristics are compensated differently at different firms (Abowd et al. 1999).

The approach I take to understanding the impact of unobserved heterogeneity of young firms on the compensation of employees is to focus on measures of firm success that are observable ex post. I define success in two ways. The first is the firm's survival over a 1-, 2-, or 3-year horizon, defined as continued appearance in the dataset. One shortcoming of this definition is that non-survival can result from one of several different forms of exit, such as a bankruptcy, a voluntary closing down of a business, or a sale to another firm. A sale could be a sign of success or the result of a last-minute effort to avoid liquidation. Non-survival may therefore mix successful and unsuccessful outcomes. The dataset does not allow me to distinguish between different forms of exit. However, one advantage of measuring success in terms of survival is that the second measure of ex post success, described below, can only be calculated conditionally on the firm appearing in the dataset in future periods.

The second way in which I define ex post success is whether the young firm is in the top quartile of the overall annual productivity distribution in 1, 2, or 3 years' time, conditional on survival. The best available measure of productivity for the entire universe of British firms captured by the BSD is real sales per employee, which serves as a proxy for the average product of labour.¹⁸ While selection effects may lead the least productive firms to exit over time, the surviving firms are nevertheless heterogeneous in terms of observed productivity based on this measure. Table 1.D3 in the Appendix shows that 56% of firms aged 3 years or less are in the bottom quartile of the overall annual productivity distribution, while only 8% are in the top quartile. It is therefore relatively uncommon for an employee to work at a young firm that is in the top quartile of the productivity distribution, though some young

¹⁸While other survey data can be used to construct other measures of labour productivity, such as gross output per employee or value added per employee, that data is only available for a sample of British firms. Small firms are not sampled every year, which makes tracking their performance difficult.

firms do attain this level of performance.

I use these two measures of success to answer three questions. First, do young firms that turn out to be successful pay higher or lower wages to their early new hires than young firms that turn out to be unsuccessful? Second, how do the wages of new hires at successful young firms compare to the wages of new hires at successful mature firms? And third, how do the wages of new hires at unsuccessful young firms compare to the wages of new hires at unsuccessful mature firms?

For starting wages, these questions can be answered by enhancing the empirical model of individual wages specified by equation (1.1), with full controls and individual fixed effects, to additionally include indicators of ex post firm success in $k = 1, 2$, or 3 years' time and interacting them with the firm age dummies. When success is measured as survival, the indicator variable $\mathbb{1}(\text{Success})_{j,t+k}$ is binary. On the other hand, when success is measured in terms of productivity conditional on survival, I use three dummy variables for quartiles $q = 1, 2$, or 3 of the distribution of sales per employee, $\mathbb{1}(\text{Qtile} = q)_{j,t+k}$.

Table 1.7 reports the results.¹⁹ In Columns (1)-(3), success is defined as survival over various time horizons. In Columns (4)-(6), a successful firm is defined as one that is in the top quartile of the overall annual distribution of productivity, while an unsuccessful firm is defined as being in the bottom quartile. The results in Panel A suggest that young firms that are successful ex post in year $t + 1$, $t + 2$, or $t + 3$ pay higher starting wages to their new hires in year t than young firms that turn out to be unsuccessful in the same future periods. In the case of surviving young firms, this premium is in the range of 1.3 to 1.9% over the 1 to 3-year

¹⁹An alternative, complementary way to link the wages of early hires to the subsequent success of the firm is to ask whether wages are a predictor of firm success. This entails swapping the dependent and independent variables in Table 1.7 and regressing success indicators on wages and controls, such that the independent variables are measured earlier than the dependent variable. The results, reported in Table 1.D4 in the Appendix, confirm that over a 3-year horizon, there is a positive and statistically significant relationship between the wages of new hires and subsequent firm survival and productivity outcomes. Wages can, therefore, be thought of as a predictor of firm success.

horizon. The results for firms that survive the next 2 or 3 years are statistically significant at a 5% and 10% level, respectively. Moreover, firms in the top quartile of productivity in 1, 2, or 3 years' time, conditional on survival, pay between 5.6% and 6.4% more to their new hires today than firms that end up in the bottom quartile. The results for productivity are larger in magnitude than the results for survival and are statistically significant at a 1% level. There are two possible interpretations of the difference in the magnitude and significance of the results for the two measures of firm success. The first is that mere survival is not a sufficiently precise measure of success, as discussed earlier. The second is that while firms that survive ex post may differ ex ante in some aspects from firms that do not survive, those aspects are less relevant for wages than the attributes that differentiate high-productivity firms from low-productivity firms.

Comparing the wages of new hires at young firms with the wages of new hires at mature firms with the same ex post level of success in Panels B and C confirms the earlier, more general result that young firms pay a 1-2% wage premium relative to mature firms. This is somewhat surprising, since the same objective measure of success is likely to be more difficult for a young firm to achieve than for a mature firm. For instance, the average young firm that breaks into the top quartile of productivity may have a particularly efficient technology, innovative intangible capital, or high-quality human capital, compared to the average mature firm in the same productivity quartile. Alternatively, if young firms are more susceptible to negative shocks or if they are less prepared to weather those shocks than mature firms, then a greater proportion of young firms that are unsuccessful ex post would have potentially had positive prospects ex ante, relative to mature firms. In these scenarios, the average ex ante wage premium at young firms relative to mature firms in the same ex post success category could differ from the one estimated earlier over the entire sample.

One possible reason why the results are similar is that young firms differ systematically from mature firms on other unobserved measures. The fact that the estimated premium is higher when firm fixed effects are used, as mentioned earlier, supports this interpretation. For instance, financial constraints driven by a short track record may be putting downward pressure on starting wages at young firms, even at those that turn out to be highly productive later.

Does firm success also lead to higher wage growth over time? Figure 1.6 plots the returns within continuing matches at young firms that are in the top half of the overall productivity distribution and compares them to the average mature firm that was plotted earlier in Panel B of Figure 1.5. It turns out that the wage-tenure profile continues to be flat. However, thanks to the high initial wage premium paid to new hires at successful young firms, joining such a firm pays more in present value over the first 5 years than joining an average mature firm, for any discount rate. From a wage standpoint it matters, therefore, what kind of young firm a given individual joins.

1.4.2 Where in the economy are the high-paying jobs at young firms?

Which sectors of the economy account for the largest proportion of jobs created by young firms? And do young firms in those sectors tend to offer high-paying and stable jobs?

Table 1.8 lists the ten 2-digit SIC sectors that account for the largest number of jobs at young firms in the first year of operations (“start-ups” for short), on average across the sample period—as well as the 10 sectors that contribute the least to the stock of jobs at newly-created firms. For each sector, the table shows the average real hourly wage in the sample (expressed in 2014 pounds), the average 3-year exit rate of start-ups, and the average real productivity of start-ups after 3 years, conditional on survival and measured as real sales per employee. Out of the 46 two-digit sectors in the sample, the top ten account for 87.8%

of jobs at newly-created firms. The largest, “Other business activities,” which represents 25.5% of such jobs, captures a mix of skilled service activities that include tax, accounting, management, and technical advisory and consulting services. The average wage in this sector is £11.57 per hour, 11.4% above the overall average of £10.39, while the exit rates and productivity outcomes are close to average.

The next two top sectors offer a different picture. “Hotels and restaurants” and “Retail trade except motor vehicles” jointly account for 28.6% of jobs created at start-ups but offer average wages that are well below the overall average (at £7.21 and £8.26 per hour, respectively) and exit rates that are above average. Young firms in the hotel and restaurant sector also attain low productivity levels, even conditional on survival. While retail traders appear to achieve high productivity, this is likely to be an artefact of how the productivity variable is constructed in the table. As the only way to measure productivity for the universe of firms in this dataset is using sales, the productivity of businesses that trade finished goods, like retailers, will be overstated relative to measures based on value added that are, arguably, more relevant for wage setting.

The remaining large contributors to job creation at start-ups represent a mix of service sectors. Only in sixth place do we see “Computer and related activities,” or software and technology firms that are colloquially referred to as “start-ups.” While new firms in this sector offer high wages, as one might expect, they represent only 5.5% of aggregate employment at firms of that age. Finally, the ten sectors that contribute the least to job creation at start-ups tend to offer above-average wages but jointly represent only 1.4% of the stock of jobs at start-ups. Overall, therefore, the table paints a mixed picture of the quality of jobs at start-ups, with a large proportion of them being concentrated in low-wage and relatively unstable sectors.

Differences in average wages across sectors may be affected by considerations such as the sorting of individuals with different characteristics among different occupations and segments of the economy. Figure 1.7 assesses the relationship between relative sectoral wages at start-ups, exit rates, and productivity outcomes—after controlling for the observable characteristics of firms and workers, as well as for individual fixed effects. It is, therefore, a way to assess the quality of jobs created at start-ups after taking into account the sorting and selection of employees across sectors.

The figure plots the average exit rates and productivity levels of firms aged 1 year against the relative wages of new hires at such firms, by sector. In Panel A, the vertical axis shows the 3-year exit rate of firms aged 1 year, calculated annually for each 3-digit SIC sector and averaged over the period from 2002 to 2012. In Panel B, the vertical axis shows the sectoral average of the log of real sales per employee 3 years after firm birth, conditional on appearance in the dataset. On the horizontal axis in both panels, average sectoral wages are calculated as sector fixed effects from estimating equation (1.1) with full controls on the sub-sample of employees at firms aged 1 year. Each bubble represents a 3-digit SIC sector, and its size is proportional to the annual number of jobs at 1-year-old firms in that sector, averaged over the period from 2002 to 2012. Selected sectors with extreme values are highlighted.

Figure 1.7 confirms that jobs in sectors where young firms pay high wages, fail less frequently, and achieve high productivity, are rare compared to sectors where young firms have less attractive prospects, even after accounting for compositional differences. A substantial proportion of jobs at young firms are concentrated in a few low-wage, high-risk, and low-productivity sectors like bars, restaurants, food retail, and industrial cleaning. In contrast, the sectors with high wages, low risk, and high productivity are more heterogeneous. They

can be described as mainly high-skilled, specialised service sectors, such as software and computer services, real estate activities, architectural and engineering services, and finance. In aggregate, however, young firms in these high-wage sectors create a small proportion of jobs compared to young firms in the low-wage sectors, as indicated by the small relative size of the bubbles.

These conclusions are not surprising. Most new firms have low productivity, and their founders do not aim to innovate or grow (Shane 2009; Hurst and Pugsley 2011), so one might expect that they do not create high-paying jobs. This paper is the first, however, to confirm explicitly that variation in the quality (defined in terms of survival and profitability) of young firms is closely linked to the quality (defined in terms of wages) of the jobs that they create. It does so thanks to the availability of longitudinal data on individuals and firms. These data enable a careful definition of firm age, entry, and exit, as well as considering differences in observable and (time-invariant) unobservable characteristics of the employees of those firms, shedding some light on the characteristics of the young firms that create “good,” high-paying jobs. While the job-creating properties of young firms have captivated economists and policymakers alike, the results of this paper show that “good” jobs created by those firms are rare.

Should policy aim to stimulate firm creation in the high-wage sectors highlighted in this section to enable more high-wage jobs to be created? There are at least two reasons for approaching such a recommendation with caution. First, the relationships identified in this paper may not hold in future periods. For instance, sectors with low exit rates in the past may not be the sectors with low exit rates in the future. Second, industrial policy focused on subsidizing or otherwise encouraging firm creation in specific sectors may have implications for aggregate productivity and welfare that are not explored here. Increasing the number

of new firms in a given sector may lead to a fall in the average productivity of firms in that sector and reduce wages. The results of the previous section suggest that focusing on firms with better survival prospects and expectations of higher productivity may be more useful than a blanket sector-based policy. An important step in this direction is Guzman and Stern's (2017) methodology for forecasting the probability of future growth of early-stage start-ups based on information available in business registration records or early milestones. Future research should also focus on how to identify *ex ante* which young firms are likely to attain high productivity levels.

1.5 Declining quality of young firms and real wage stagnation

Young firms that turn out to be more successful, in terms of survival and revenue generation, pay higher wages to their new hires at the outset than young firms that are less successful. On average, however, young firms tend to have characteristics that are associated with low wages. They are small and more prevalent in sectors where wages and productivity are low, and exit rates are high. Together with the substantial tenure wage premium that accrues to many employees at large firms, these factors explain—in a mechanical sense—the striking wage differential between young and mature firms observed in the raw data on all jobs and illustrated earlier in Figure 1.1.

Perhaps even more concerning from a wage perspective is the fact that wages at young firms have worsened relative to mature firms over time. Figure 1.8 shows this for all employee jobs in the dataset. Between 2002 and 2009, the mean real hourly wage at firms aged 3 years or less declined by 4.6%. In contrast, the mean real hourly wage at firms aged 10 years or more grew by 9.4% (based on log differences). Even though average real wages at mature firms then began to fall after the 2008 financial crisis, wages at young firms declined at a faster rate. Between 2009 and 2012, the mean real wage at mature firms declined by 8.3%,

while the mean real wage at young firms declined by 11.4%. Although the overall real wage decline in the UK is a well-documented phenomenon that has generated a growing body of research (Gregg et al., 2014), the relative worsening of wages at young firms that I provide evidence for here is a new result that, to the best of my knowledge, has not been documented in the literature.

The relative worsening of wages at young firms has contributed substantially to the aggregate real wage decline in the UK over the recent decade, since young firms account for a sizeable share of job creation in the economy. This is the conclusion of a simple counterfactual scenario presented in Table 1.9, which shows that the widening of wage differences across broad firm age categories can account for 60% of the fall in average real wages in the UK since the financial crisis. For simplicity, the sample period is split into two periods of equal length, 2002-2008 and 2009-2015, which attenuates the peak-to-trough fluctuations in wages. Firms are divided into four age categories: births, 2 to 3 years old, 4 to 9 years old, and aged 10 or more. The table shows that average real wages declined within all age categories between the first and second period, and aggregate real wages fell by 4.6%. However, wages declined more at young firms than at mature firms. Had the relative proportional differences in wages across the four age categories remained constant, the aggregate wage decline would have been “only” 1.7%, an improvement of 2.6 percentage points (or 60%). This result holds despite the fact that mature firms’ share of total employment increased slightly between the two periods. This simple thought experiment highlights the relevance of wages at young firms to the aggregate economy.

To understand whether the relative worsening of wages at young firms compared to mature firms is due to job tenure, compositional effects, or selection effects, I follow similar steps as in the earlier discussion of the average wage difference in Section 1.3.2. Figure 1.9

focuses on the average wage premium (or penalty) at firms aged 3 years or less, relative to firms aged 10 years or more. It begins by comparing wages in all jobs at young and mature firms. It then looks only at the wages of new hires to assess whether the impact of job tenure on the wage differential between young and mature firms has changed over time. Full controls are then added to check whether the trend still holds in the sub-sample of new hires after controlling for compositional effects. Finally, individual fixed effects are added and the sample expanded to all employees in the dataset to explore the role of individual selection.

The widening gap in wages between young and mature firms observed in the raw data is due largely to the changing composition of young firms and the jobs that they offer. Panel A in Figure 1.9 compares all jobs and shows the widening gap in wages between the two firm age categories highlighted in the previous figure. The horizontal dashed line marks the average wage gap of -12.6% (log difference -0.135) that was estimated earlier via regression in the first bar of the bar chart in Figure 1.3. Over time, however, the wage gap worsened from -2.1% in 2002 to -19.2% (log difference -0.213) in 2015. Moreover, Panel B shows that the downward trend in relative wages at young firms is present even for new hires. Between 2002 and 2015, the wage difference between young and mature firms changed from 12.7% (log difference 0.119) to -3.5%. Thus, the slight average premium estimated earlier, when pooling all data across time, has turned into a penalty in recent years (although the estimated wage premium is significantly below zero only in 2015). However, the slope of the general trend is similar to that in Panel A, suggesting that the effect of job tenure on the wage differential between young and mature firms has remained roughly constant across time. This confirms earlier conclusions drawn from pooled data. In contrast, Panel C shows that the trend flattens substantially when the annual wage differential is estimated

with full controls. The estimated wage gap now changes from a premium of 1.7% in 2002 to a penalty of -1.8% in 2015. While the trend is still negative, it is far less steep than before controlling for observables. This flattening of the trend suggests that there has been a relative worsening in the observable quality of either the young firms themselves or the individuals that they hire, compared to mature firms.

The main explanation for why the wages of new hires at young firms have declined so substantially relative to the wages of new hires at mature firms over the entire sample period is the shift in the sectoral composition of young firms and the occupational structure of the jobs that they offer. Figure 1.10 illustrates the changing sectoral composition. It shows the annual distribution of jobs at young and mature firms in the matched sample among four large sectors from 2002 to 2015. The share of jobs at young firms has shifted significantly more towards low-wage sectors like Accommodation and Restaurants, compared to mature firms over the same period. At the same time, the share of jobs in Trade and Other Non-Financial Services at young firms has grown less than at mature firms, while the decline in Manufacturing and Construction jobs has been steeper. Changes in the occupational structure of jobs at young and mature firms are illustrated separately in Figure 1.11. Using four broad occupational skill categories following ONS (2003), this figure shows that jobs at young firms have become relatively more concentrated in the two lowest-skilled categories, 1 and 2, than jobs at mature firms over the same period. These shifts in sectoral and occupational composition help explain the divergence in the wages of new hires at young firms, relative to mature firms, observed in Panel B of Figure 1.9.

Average wages at young firms have declined relative to wages at mature firms since 2012, even after accounting for individual fixed effects. In Panel D of Figure 1.9, a wage premium of 1.8% at young firms in 2012 turned into a wage penalty of 2.8% in 2015, both statistically

significant at the 5% level. Since these estimates control for unobservable (time-invariant) individual characteristics, there are at least two possible explanations for this decline. The first is that the characteristics of the employees hired by young firms in the recent period have changed in some unobservable way relative to mature firms over the same timeframe. Another explanation is that the nature of young firms has changed in some unobserved way that differs from the way that mature firms have changed during the same period.²⁰

In what follows, I focus on the latter explanation by asking whether the prospects of young firms have worsened since 2012. To do so, I consider the measures of firm success discussed earlier: survival and productivity, measured as real sales per employee. Figure 1.12 shows that the one-year exit rates of new entrants increased by 4 percentage points between 2012 to 2014, compared to mature firms in the same sectors and size categories (the total one-year exit rate of entrants increased from 15.9% to 20.2% over that period). The exit rates observed at the end of the sample period are greater than those observed during any other part of the business cycle since 2002. One reason for caution, however, is that these estimates could be affected by the proximity to the end of the sample period, such that there are limited opportunities to observe the firm again in the future. For instance, they might be driven to some extent by firms that go dormant instead of exiting permanently. However, the conclusion about the worsening prospects of young firms is also confirmed by Figure 1.13, which documents the falling productivity of young firms relative to mature firms in the same sectors and size categories over the same period.

To confirm whether the worsening prospects of young firms are a relevant factor in explaining the decline in wages relative to mature firms, I modify the regressions with full controls and individual fixed effects to account for ex post firm success. Figure 1.14 shows

²⁰Appendix B shows that the dip in relative wages is unlikely to be driven by potential firm owners, self-employment, or the restructuring of employment relationships into subcontracting relationships. Testing these explanations is important, because self-employment has grown substantially in the UK in recent years.

that the trend in the wage difference between young and mature firms is now flatter and does not exhibit the sharp drop after 2012 when I consider only successful firms. Panel A plots this wage difference by year, using only to the sub-sample of firms that survive the following year. The time trend is slightly negative, but the post-2012 dip is far less pronounced than before, when both surviving and non-surviving firms were included. Panel B plots the wage difference for the sub-sample of firms with above-median productivity one year forward, conditional on survival. The average wage premium at young firms relative to mature firms is higher over time than in the sample of surviving firms. Although the standard errors are wide, due to a small number of young firms that achieve high productivity levels in any given year, the post-2012 dip in relative wages at young firms disappears completely. Thus, an individual who joins a highly productive young firm can still be expected to earn similar wages to those at a highly productive mature firm.

In all, the patterns described in this section highlight a substantial divergence between the recent trends in average wages at young and mature firms. A high-level thought experiment suggests that the widening wage gap between young and mature firms accounted for more than half of the fall in average real wages in the UK between the pre- and post-crisis periods. This trend coincides, more recently, with a reduction in the survival rate and productivity of young firms. The close historical link between wages and ex post success of young firms suggests that the recent widening of the wage gap relative to mature firms augurs weak future performance of the average start-up. Given the significance of young firms to job creation and other aggregate economic outcomes, this may be a cause for concern.

1.6 Conclusions

In this paper, I use matched employer-employee data for Great Britain to study the wages of employees at young firms. I document and explain several novel findings. I show that most of the differences in the distributions of hourly wages, weekly earnings, and hours between young and mature firms observed in the raw data on all jobs are driven by tenure effects. New hires at young firms actually earn a small premium relative to similar new hires at mature firms. This result is robust to using several approaches to addressing the selection of individuals into young firms based on observable and time-invariant unobservable characteristics. I also find that despite the initial premium at young firms, new hires experience better subsequent earnings prospects at mature firms.

Young firms are a heterogeneous group, however. Those young firms that survive and attain high productivity levels, conditional on survival, pay higher wages to new hires at the outset than less successful young firms. These results suggest that policies aiming to stimulate job creation by means of new firm creation may have different labour market effects, depending on the types of businesses that form as a result.

The results of this paper give rise to a number of questions for future research. For instance, why do people join a typical young firm, despite it offering only a small initial wage premium and worse earnings growth than a typical mature firm? There are at least three possible explanations. First, the individuals who join young firms may be less risk-averse than those who join mature firms (Roach and Sauermann 2017). Second, employees may be attracted to young firms for non-pecuniary motives, such as working in a more flexible environment (Koch et al. 2013) or a flatter hierarchy (Sørensen 2007). This is analogous to the argument that non-pecuniary reasons may explain why many individuals decide to become entrepreneurs despite the average returns being lower than in the counterfactual

scenario of wage employment (Hamilton 2000; Hurst and Pugsley 2011, 2017). Third, people may join young firms simply because other, preferred forms of employment are difficult to obtain in some labour markets, particularly in weak economic conditions. All these explanations merit further exploration, using the results of this paper as a starting point.

Continued research on the employees of young firms and their earnings prospects is vital, given the quantitative importance of entrepreneurship for aggregate job creation and the fact that the nature of young firms appears to be changing. Over the last decade and a half, jobs at young firms have become more likely to be concentrated in low-wage sectors and occupations. Young firms have also become less productive—both in absolute terms and relative to mature firms operating within the same narrowly-defined sector. Moreover, these trends have coincided with a deterioration in the wage gap between young and mature firms. Highly-paid and stable jobs at young firms have thus become increasingly rare, casting a shadow on the contribution of young firms to job creation. From a labour market perspective, the inevitable conclusion is that indiscriminate support for new firm creation may be detrimental to the quality of new jobs. Any local or national policies that aim to stimulate job growth through the creation of new businesses may do well to pay close attention to the nature of the firms that are being created.

1.7 Bibliography

- Abowd, John, Francis Kramarz, and David Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67(2): 251-333.
- Adrian, Pawel and Brian Bell. 2018. "Pension Shocks and Wages." CEP Discussion Paper No. 1536.
- Anyadike-Danes, Michael, Mark Hart, and Jun Du. 2015. "Firm Dynamics and Job Creation in the United Kingdom: 1998-2013." *International Small Business Journal* 33(1): 12-27.
- Becker, Randy, John Haltiwanger, Ron Jarmin, Shawn Klimek, and Daniel Wilson. 2006. "Micro and Macro Data Integration: The Case of Capital," in D. Jorgenson, S. Landefeld, and W. Nordhaus (eds.), *A New Architecture for the U.S. National Accounts*, Chicago: NBER/University of Chicago Press.
- Bell, Brian and Stephen Machin. 2018. "Minimum Wages and Firm Value." *Journal of Labor Economics* 36(1): 159-195.
- Brixy, Udo, Susanne Kohaut, and Claus Schnabel. 2007. "Do Newly Founded Firms Pay Lower Wages? First Evidence from Germany." *Small Business Economics* 29(1): 161-171.
- Brown, Charles and James Medoff. 2003. "Firm Age and Wages." *Journal of Labor Economics* 21(3): 677-697.
- Burton, Diane, Michael Dahl, and Olav Sorenson. 2017. "Do Startups Create Good Jobs?" *Industrial and Labor Relations Review*, forthcoming.
- Couch, Kenneth and Dana Placzek. 2010. "Earnings Losses of Displaced Workers Revisited." *American Economic Review* 100(1): 572-589.
- Cox, James and Ronald Oaxaca. 1989. "Laboratory Experiments with a Finite-Horizon Job Search Model." *Journal of Risk and Uncertainty* 2(3): 301-329.
- Criscuolo, Chiara, Peter Gal, and Carlo Menon. 2014. "The Dynamics of Employment Growth: New Evidence from 18 Countries." OECD Science, Technology and Industry Policy Papers No. 14. Paris: OECD Publishing.
- Dahl, Michael and Steven Klepper. 2015. "Whom Do New Firms Hire?" *Industrial and*

Corporate Change 24(4): 819-836.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. "The Role of Entrepreneurship in US Job Creation and Economic Dynamism." *Journal of Economic Perspectives* 28(3): 3-24.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2016. "Where Has All the Skewness Gone? The Decline in High-Growth (Young) Firms in the U.S." *European Economic Review* 86(1): 4-23.

Evans, David and Boyan Jovanovic. 1989. "An Estimated Model of Entrepreneurial Choice Under Liquidity Constraints." *Journal of Political Economy* 97(4): 808-827.

Gregg, Paul, Stephen Machin, and Mariña Fernández-Salgado. 2014. "Real Wages and Unemployment in the Big Squeeze." *Economic Journal* 124(576): 408-432.

Guzman, Jorge and Scott Stern. 2017. "Nowcasting and Placecasting Entrepreneurial Quality and Performance," in J. Haltiwanger, E. Hurst, J. Miranda, and A. Schoar (eds.), *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, Chicago: University of Chicago Press.

Haltiwanger, John, Ron Jarmin, and Javier Miranda. 2013. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95 (2): 347-361.

Hamilton, Barton. 2000. "Does Entrepreneurship Pay? An Empirical Analysis of the Returns of Self-Employment." *Journal of Political Economy* 108(3): 604-631.

Heyman, Fredrik. 2007. "Firm Size or Firm Age? The Effect on Wages Using Matched Employer-Employee Data." *Labour* 21(2): 237-263.

Hurst, Erik and Benjamin Pugsley. 2011. "What Do Small Businesses Do?" *Brookings Papers on Economic Activity*, 73-142.

Hurst, Erik and Benjamin Pugsley. 2017. "Wealth, Tastes, and Entrepreneurial Choice," in J. Haltiwanger, E. Hurst, J. Miranda, and A. Schoar (eds.), *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, Chicago: University of Chicago Press.

Huttunen, Kristiina, Jarle Møen, and Kjell Salvanes. 2011. "How Destructive Is Creative Destruction? Effects of Job Loss on Job Mobility, Withdrawal and Income." *Journal of the European Economic Association* 9(5): 840-870.

- International Entrepreneur Rule. 2016. *Federal Register* 81 FR 60129.
- Jolivet, Grégory, Fabien Postel-Vinay, and Jean-Marc Robin. “The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the US.” *European Economic Review* 50(4): 877-907.
- Jovanovic, Bojan. 1982. “Selection and Evolution of Industry.” *Econometrica* 50(3): 649-670.
- Kambourov, Gueorgui and Iourii Manovskii. 2009. “Occupational Specificity of Human Capital.” *International Economic Review* 50(1): 63-115.
- Koch, Andreas, Daniel Pastuh, and Jochen Späth. 2013. “New Firms and New Forms of Work.” *mimeo*.
- Manso, Gustavo. 2016. “Experimentation and Returns to Entrepreneurship.” *Review of Financial Studies* 29(9): 2319-2340.
- Michelacci, Claudio and Olmo Silva. 2007. “Why So Many Local Entrepreneurs?” *Review of Economics and Statistics* 89(4): 615-633.
- Neal, Derek. 1995. “Industry-Specific Human Capital: Evidence from Displaced Workers.” *Journal of Labor Economics* 13(4): 653-677.
- Nosal, Ed and Peter Rupert. 2007. “How Amenities Affect Job and Wage Choices over the Life Cycle.” *Review of Economic Dynamics* 10(3): 424-443.
- Nyström, Kristina and Gulzat Zhetibaeva Elvung. 2015. “New Firms as Employers: The Wage Penalty for Voluntary and Involuntary Job Switchers.” *Labour* 29(4): 348-366.
- Nyström, Kristina and Gulzat Zhetibaeva Elvung. 2014. “New Firms and Labor Market Entrants: Is There a Wage Penalty for Employment in New Firms?” *Small Business Economics* 43(2): 399-410.
- OECD. 2016. *Entrepreneurship at a Glance 2016*, OECD Publishing, Paris.
- Office for National Statistics. 2006. *Business Structure Database User Guide*. <https://discover.ukdataservice.ac.uk/catalogue?sn=6697> (last accessed 16 March 2017).
- Office for National Statistics. 2017. Annual Survey of Hours and Earnings, 1997-2016: Secure Access. [data collection]. 9th Edition. UK Data Service. SN: 6689, <http://doi.org/10.5255/UKDA-SN-6689-8> (last accessed 16 March 2017).

- Office for National Statistics. 2017. Business Structure Database, 1997-2016: Secure Access. [data collection]. 7th Edition. UK Data Service. SN: 6697, <http://doi.org/10.5255/UKDA-SN-6697-7> (last accessed 16 March 2017).
- Ouimet, Paige and Rebecca Zarutskie. 2014. "Who Works for Startups? The Relation between Firm Age, Employee Age, and Growth." *Journal of Financial Economics* 112(3): 386-407.
- Roach, Michael and Henry Sauermann. 2017. "Ex Ante Career Preferences and Sorting into Startup Employment." *mimeo*.
- Shane, Scott. 2009. "Why Encouraging More People to Become Entrepreneurs Is Bad Public Policy." *Small Business Economics* 33(2):141-149.
- Small Business Administration. 2014. *Strategic Plan. Fiscal Years 2014-2018*. https://www.sba.gov/sites/default/files/aboutsbaarticle/SBA_FY_2014_-2018_Strategic_Plan-1.pdf (last accessed 30 September 2017).
- Sørensen, Jesper. 2007. "Bureaucracy and Entrepreneurship: Workplace Effects on Entrepreneurial Entry." *Administrative Science Quarterly* 52(3): 387-412.
- Sullivan, Paul. 2010. "Empirical Evidence on Occupation and Industry-Specific Human Capital." *Labour Economics* 17(3): 567-580.
- US Census Bureau. 2015. Business Dynamics Statistics. [data collection]. https://www.census.gov/ces/dataproducts/bds/data_firm.html (last accessed 30 September 2017).
- Zangelidis, Alexandros. 2008. "Occupational and Industry Specificity of Human Capital in the British Labour Market." *Scottish Journal of Political Economy* 55(4): 420-443.

1.8 Figures and tables

Figure 1.1: Average wage gap between young and mature firms

This figure plots firm age coefficients from a regression of (a) log real weekly earnings and (b) log real hourly wages on firm age dummies and year fixed effects. The omitted category are firms aged more than 10 years. Real earning and wages are deflated by the Consumer Price Index and expressed in 2014 GBP. Vertical bars indicate 95% confidence intervals based on clustering by firm. The sample consists of 1,329,517 observations pooled.





Figure 1.2: Distribution of hourly wages by firm age

This figure compares the relative distribution of real hourly wages at firms aged 3 years or less (shaded bars) with wages at firms aged 10 years or more (transparent bars). Data is pooled over all years in the sample (2002-2015). Wages are deflated by CPI and expressed in 2014 GBP. Panel A shows all jobs, while panel B shows new hires.

Panel A: All Employees



Panel B: New Hires Only




Figure 1.3: New hire wage premium at young vs. mature firms

This figure plots the estimated coefficient on the indicator for a firm aged 3 years or less from individual wage regressions with varying controls, specified by equation (1.1). Detailed results are shown in Columns (1)-(5) of Table 1.3.




Figure 1.4: New hire wage premium at young vs. mature firms (selection on unobservables)

This figure plots the estimated coefficient on the indicator for a firm aged 3 years or less from individual wage regressions on various sub-samples that represent four approaches to address potential selection of individuals to young and mature firms based on unobservables. The models are specified by equation (1.1). Detailed results are shown in Columns (1)-(4) of Table 1.4.





Figure 1.5: Returns to joining a young firm

This figure presents the evolution of wages of individuals hired by young firm, compared to individuals hired by a mature firm. Panel A shows the subsequent evolution of wages regardless of whether the individuals stay at the same firm or are observed at other firms (unconditional returns). Panel B shows returns to tenure within continuing matches. It plots the coefficients on tenure dummies from a regression of individual wages on tenure dummies interacted with firm age dummies, full controls, and individual fixed effects. In each plot, wages are expressed relative to new hire wages at mature firms. Thin grey lines represent 95% confidence intervals.

Panel A: Unconditional Returns



Panel B: Returns in Continuing Matches




Figure 1.6: Returns to joining a successful young firm

This figure presents the returns to tenure within continuing matches of individuals at young firms with above-median productivity, compared to individuals at any mature firm. In each plot, wages are expressed relative to new hire wages at mature firms. Thin grey lines represent 95% confidence intervals. See the notes to Figure 1.5 for more details.





Figure 1.7: Where are the “good” jobs at young firms?

This figure plots sectoral measures of the riskiness and productivity of young firms against the average wages of new hires at those firms. In both panels, average wages of new hires are expressed on the horizontal axis as the estimated sector fixed effects from a regression of individual log real hourly wages of new hires at firms aged 1 year on 3-digit sector dummies and full controls from Column (5) of Table 1.3 (7,565 observations). In panel A, the vertical axis shows the sectoral average of the 3-year exit rate of firms aged 1. In panel B, the vertical axis is the sectoral average of log real sales per employee 3 years after firm birth, conditional on survival. These averages are calculated over the period from 2002 to 2012, using BSD data. The size of each bubble represents the average employment of firms aged 1 year in the sector. To comply with disclosure restrictions, only sectors with at least 10 new hires observed over the sample period are shown.

Panel A: Sector wages vs. exit rates of young firms



Panel B: Sector wages vs. productivity of young firms




Figure 1.8: Wages by firm age over time

This figure plots the mean log real hourly wage by firm age category from 2002 to 2015. Wages are deflated by the Consumer Price Index and expressed in 2014 pounds sterling (GBP). Light grey solid lines represent 95% confidence intervals. The total number of observations is 1,322,716.




Figure 1.9: Young firm wage premium/penalty by year

This figure shows the average wage premium at firms aged 3 years or less relative to firms aged 10 years or more. The vertical axis represents the marginal effect of working for a young firm by year, based on individual regressions of log real hourly wages on firm age dummies interacted with year dummies. Panel (a) compares all jobs. Panel (b) looks at the sub-sample of new hires. Panel (c) adds full controls, as described in the main text. Panel (d) adds individual fixed effects.




Figure 1.10: Shifts in the sectoral composition of young firms

This figure compares the trends in the sectoral share of jobs at young firms aged 3 years or less vs. mature firms aged 10 years or more. The sectors shown are: Manufacturing and Construction (sectors D and F in the SIC 2003 classification), Trade (G), Accommodation and Restaurants (F), and Other Non-Financial Services (K). Finance (J) and Transportation (I) are not shown due to a small sample of young firms in those sectors. The charts show that young firms have become more concentrated in Accommodation and Restaurants over this period, both in absolute terms and relative to mature firms over the same period.




Figure 1.11: Shifts in the occupational composition of jobs at young firms

This figure compares the trends in the proportion of jobs within 4 occupational skill categories at young firms aged 3 years or less vs. mature firms aged 10 years or more. The occupational skill categories are based on the ONS classification. The vertical dashed line in each plot represents a change in the occupational classification system from SOC 2000 to SOC 2010, implemented in the dataset from 2012 onwards, which leads to a break in the series. The charts show that the share of low-skilled jobs at young firms has grown faster than the share of low-skilled jobs at mature firms.




Figure 1.12: Rising exit rates of young firms

This figure plots year fixed effects from regressions of a 0/1 indicator of exit in the following year on year, 5-digit sector, and firm size dummies, separately for firms aged 1 year and 10 or more years, based on the BSD. The baseline (omitted) year is 2002. Firm size dummies consist of 19 ventiles of employment. The chart shows that while the trends in the exit rate of young and mature firms began to diverge slightly in the mid-2000s, there was an additional significant increase in relative exit rates of new entrants, compared to mature firms aged 10 or more, after 2012. The total number of observations of exiting firms in these two age categories is 1,240,330, and the differences between the mean exit rates of firms in the two age categories are highly statistically significant in all years.




Figure 1.13: Falling productivity of young firms

This figure plots year fixed effects from regressions of log real sales per employee on year, 5-digit sector, and firm size dummies, separately for firms aged 1, 2, and 10 or more years, based on the BSD. The baseline (omitted) year is 2002. Firm size dummies consist of 19 ventiles of employment. The chart shows that while the within-sector productivity trends of young and mature firms began to diverge in the mid-2000s, there was an additional significant drop in relative productivity of new entrants, compared to mature firms aged 10 or more, in 2012. The total number of observations is 24,810,907, and the differences between the mean productivity levels of firms in different age categories are highly statistically significant in most years.




Figure 1.14: Young firm wage premium/penalty by year and ex post firm success

This figure shows the average wage premium at firms aged 3 years or less relative to firms aged 10 years or more. The vertical axis represents the marginal effect of working for a young firm by year for a new hire, based on individual regressions of log hourly wages on firm age dummies interacted with year dummies. All regressions include full controls, as described in the main text, and individual fixed effects.



Table 1.1: Summary statistics: young firms

This table presents summary statistics for observations at firms aged 3 years or less and at firms aged 4-9 years. Real variables are deflated using CPI and expressed in 2014 GBP. Occupational skill levels follow ONS (2003) and are derived from 2-digit SOC 2000 and 2010 occupational classifications, based on the duration of training and/or work experience typically required to perform the activities in a given job. Sectoral classifications are based on the SIC 2003 system, with Industry defined as sectors C-E, Construction as sector F, and Services as sectors G-K.

	Observations at firms aged ≤ 3 years				Observations at firms aged 4-9 years			
	Mean	Median	Standard Deviation	Number of Observations	Mean	Median	Standard Deviation	Number of Observations
<i>Individual characteristics</i>								
Real hourly wage (£)	12.9	9.7	11.2	41,132	14.2	10.7	12.6	88,780
Log(Real hourly wage)	2.4	2.3	0.5	41,132	2.5	2.4	0.5	88,780
Real weekly earnings (£)	436.9	343.1	431.3	41,132	497.4	396.1	470.7	88,780
Weekly hours	32.5	37.0	13.3	41,132	34.1	37.5	12.6	88,780
Age	38.2	37.0	12.7	41,132	38.8	38.0	12.6	88,780
Female (0/1)	0.4	0.0	0.5	41,132	0.4	0.0	0.5	88,780
Job tenure	1.5	1.0	0.9	38,494	3.3	3.0	2.3	84,939
<i>Firm characteristics</i>								
Employment	154.8	9.0	743.3	41,132	298.8	20.0	966.2	88,780
Real sales (£m)	10.7	0.5	63.2	41,132	36.9	1.6	196.2	88,780
Firm age (years)	2.1	2.0	0.9	41,132	6.4	6.0	1.7	88,780
<i>% of observations by:</i>								
	1 (Low)	2-3	4 (High)		1 (Low)	2-3	4 (High)	
Occupational skill level (%)	19.2	57.6	23.1		17.1	58.5	24.4	
	Industry	Services	Construction		Industry	Services	Construction	
Sector (%)	11.1	80.9	8.0		12.3	79.5	8.2	

Table 1.2: Summary statistics: mature firms

This table presents summary statistics for observations at firms aged 10 years or more and at firms classified as ‘not truly young,’ as defined in Section 1.2.4. See the notes to Table 1.1 for more details.

	Observations at firms aged 10+ years				Observations at ‘not truly young’ firms			
	Mean	Median	Standard Deviation	Number of Observations	Mean	Median	Standard Deviation	Number of Observations
<i>Individual characteristics</i>								
Real hourly wage (£)	14.9	11.2	14.7	1,165,860	15.8	12.1	13.3	33,741
Log(Real hourly wage)	2.5	2.4	0.6	1,165,860	2.6	2.5	0.6	33,741
Real weekly earnings (£)	538.6	435.8	493.6	1,165,860	580.1	471.2	505.1	33,741
Weekly hours	35.2	37.5	11.3	1,165,860	35.9	37.5	11.0	33,741
Age	39.9	40.0	12.8	1,165,860	39.2	38.0	12.2	33,741
Female (0/1)	0.4	0.0	0.5	1,165,860	0.4	0.0	0.5	33,741
Job tenure	8.0	5.0	8.3	1,117,167	6.4	4.0	7.7	29,967
<i>Firm characteristics</i>								
Employment	23,417.2	1,528.0	52,328.8	1,165,860	4,399.0	149.0	11,715.2	33,741
Real sales (£m)	4,282.2	162.9	11,369.1	1,165,860	855.8	9.9	2,533.9	33,741
Firm age (years)	28.1	29.0	9.2	1,165,860	6.7	7.0	2.3	33,741
<i>% of observations by:</i>								
	1 (Low)	2-3	4 (High)		1 (Low)	2-3	4 (High)	
Occupational skill level (%)	16.3	62.7	21.0		12.7	61.6	25.8	
	Industry	Services	Construction		Industry	Services	Construction	
Sector (%)	21.2	73.8	5.0		17.1	71.8	11.2	

Table 1.3: Comparing hourly wages at young and mature firms

This table reports the coefficients on firm age dummies from regressions of the log real hourly wage given by equation (1.1). Column (1) includes year fixed effects. Column (2) uses the sub-sample of new hires. Column (3) additionally controls for firm size and 5-digit sector. Column (4) adds controls for 2-digit occupation, indicators for full-time/part-time and permanent/temporary jobs, and NUTS 1 region dummies. Column (5) adds age, age squared, gender, an indicator for past young firm experience, and years of occupation-specific experience (“full controls”). Standard errors clustered at the firm level are in parentheses.

Dependent variable: log real hourly wage	Raw data	New hires	Size and sector	Job attributes	Personal attributes (“full controls”)
	(1)	(2)	(3)	(4)	(5)
Firm aged ≤ 3 years	−0.135*** (0.015)	0.025** (0.012)	0.031*** (0.005)	0.008** (0.004)	0.000 (0.004)
Firm aged 4-9 years	−0.053*** (0.016)	0.026** (0.012)	0.003 (0.005)	0.002 (0.003)	0.009** (0.003)
Not truly young	0.051* (0.027)	0.085*** (0.027)	0.025** (0.011)	0.018*** (0.007)	0.021*** (0.007)
R ²	0.01	0.01	0.31	0.59	0.70
Number of Observations	1,329,513	272,969	272,705	272,570	179,812

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.4: Four approaches to addressing potential selection

This table reports the coefficients on firm age dummies from regressions of the log real hourly wage, using four approaches to account for sorting and selection of individuals into firms of different ages based on unobservable attributes, as described in Section 1.3.3 of the main text. All regressions include full controls, as listed in the notes to Table 1.3. Standard errors clustered at the firm level are in parentheses.

Dependent variable: log real hourly wage	Individual selection on unobservables			
	Displaced workers	Labor market entrants	Individual fixed effects (labor market entrants)	Individual fixed effects (all workers)
	(1)	(2)	(3)	(4)
Firm aged ≤ 3 years	0.021 (0.020)	0.010* (0.006)	0.020*** (0.007)	0.014*** (0.003)
Firm aged 4-9 years	-0.008 (0.024)	0.002 (0.005)	-0.006 (0.006)	0.011*** (0.002)
Not truly young	0.013 (0.039)	0.003 (0.010)	0.006 (0.008)	0.006 (0.005)
R ²	0.74	0.51	0.42	0.20
Number of Observations	2,208	56,089	165,428	1,226,228

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.5: Do young firms compensate employees for high exit risk? New hire wage premium

This table reports the results of individual-level regressions of the log real hourly wage on firm age dummies as in equation (1.1), with the addition of the exit rate differential and its interaction with the firm age dummies. The exit rate differential is calculated as the difference in the average 3-year exit rates between young and mature firms within each 5-digit SIC sector over the period from 2002 to 2012. The regression includes individual fixed effects and full controls, as described in the main text. Standard errors are clustered at the firm level.

	(1)
Exit rate differential	0.086** (0.041)
Firm aged ≤ 3 years	-0.038*** (0.007)
Exit rate differential * Firm aged ≤ 3 years	0.807** (0.088)
Number of Observations	1,226,228

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.6: Do young firms compensate employees for high exit risk? Wage returns

This table reports the results of individual-level regressions of the log real hourly wage using equation (1.3) in the main text, focusing on continuing firm-worker matches. The values in the table represent the wage premium at young firms relative to mature firms based on the number of years elapsed since joining the firm (up to 4 years shown). Column 1 shows the results for the sub-sample of sectors where the exit rate differential is greater than the exit rate differential of the median sector. Column 2 uses the sub-sample of remaining sectors. The exit rate differential is calculated as the difference in the average 3-year exit rates between young and mature firms within a 5-digit SIC sector over the period from 2002 to 2012. Both regressions include individual fixed effects and full controls, as described in the main text. Standard errors are clustered at the firm level.

Year since hiring	Young firm wage premium	
	Exit rate differential > median sector (1)	Exit rate differential ≤ median sector (2)
0	0.031*** (0.005)	0.006 (0.004)
1	−0.007 (0.005)	−0.003 (0.004)
2	−0.006 (0.006)	−0.014*** (0.005)
3	−0.019*** (0.006)	−0.014*** (0.005)
4	−0.016** (0.007)	−0.022*** (0.006)
Discount rate required for NPV=0	17%	119%
Number of Observations	492,132	709,363

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.7: Wages and ex post success of young firms

This table measures the relationship between a young firm's ex-post success over a 1-, 2-, and 3-year horizon and the wages it pays to new hires. In Columns (1)-(3), success is defined as survival. In Columns (4)-(6), a "successful" ("unsuccessful") firm is defined as one that is in the top (bottom) quartile, conditional on survival. Panel A compares wages at successful young firms to wages at unsuccessful young firms. Panel B compares wages at successful young firms to wages at successful mature firms, while Panel C compares wages at unsuccessful young firms to wages at unsuccessful mature firms. All regressions include full controls, as listed in the notes to Table 1.3. Standard errors clustered by firm are in parentheses.

	Measure of firm success					
	Survival			Top quartile of productivity conditional on survival		
	1-year horizon (1)	2-year horizon (2)	3-year horizon (3)	1-year horizon (4)	2-year horizon (5)	3-year horizon (6)
Panel A: Successful young firms vs. unsuccessful young firms						
Wage premium/(penalty)	0.013 (0.012)	0.019** (0.008)	0.013* (0.007)	0.064*** (0.010)	0.063*** (0.010)	0.056*** (0.010)
Panel B: Successful young firms vs. successful mature firms						
Wage premium/(penalty)	0.015*** (0.003)	0.014*** (0.003)	0.013*** (0.003)	0.025*** (0.010)	0.020** (0.009)	0.013 (0.009)
Panel C: Unsuccessful young firms vs. unsuccessful mature firms						
Wage premium/(penalty)	0.015 (0.015)	0.001 (0.010)	0.001 (0.008)	0.017*** (0.004)	0.014*** (0.004)	0.011** (0.005)
Number of Observations	1,167,661	1,063,910	965,404	1,167,661	1,063,910	965,404

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.8: Start-up job creation, wages, and firm outcomes by sector

This table shows the top and bottom 10 sectors by contribution to the average annual number of jobs created at all firms aged 1 year (“start-ups”) between 2002 and 2015, based on firm-level data from BSD. The average real hourly wage for each sector is based on individual-level data from ASHE between 2002 and 2015 and calculated as $\exp(\text{mean of the log real hourly wage})$. The average 3-year exit rate and average productivity 3 years from entry conditional on survival are calculated based on the BSD. Productivity is measured as real sales per employee. Sectors with fewer than 10 individuals observed in ASHE as new hires at start-ups over the sample period are excluded. Real values are deflated by CPI and expressed in 2014 £.

Sector (2-digit SIC 2003 code)	Average contribution to job creation at startups (%)	Average real hourly wage (£)	Average 3-year exit rate of startups (%)	Average real productivity of startups after 3 years (£000s)
<i>Top 10 sectors by contribution to job creation at startups:</i>				
Other business activities (74)	25.5	11.57	39.0	41.6
Hotels and restaurants (55)	17.8	7.21	50.2	34.5
Retail trade except motor vehicles (52)	10.8	8.26	42.5	55.2
Construction (45)	10.5	11.18	39.3	49.5
Real estate activities (70)	6.6	12.28	29.2	39.2
Computer and related activities (72)	5.5	15.34	39.0	41.0
Wholesale trade except motor vehicles (51)	4.5	11.52	43.3	67.0
Sale, maintenance and repair of motor vehicles (50)	3.1	9.58	40.8	55.7
Land transport (60)	2.1	10.51	43.8	40.8
Publishing, printing, and reproduction of recorded media (22)	1.4	12.83	39.3	40.0
Total contribution of the top 10 sectors	87.8			
<i>Bottom 10 sectors by contribution to job creation at startups:</i>				
Research and development (73)	0.2	15.51	31.9	34.7
Manufacture of medical/optical instruments and clocks (33)	0.2	13.46	41.1	47.2
Manufacturing of pulp, paper, and paper products (21)	0.2	12.30	41.4	50.4
Manufacturing of basic metals (27)	0.2	12.10	39.0	53.5
Manufacturing of radio, TV & communication equipment (32)	0.2	12.25	42.0	46.1
Recycling (37)	0.2	12.57	43.6	68.0
Electricity, gas, steam and hot water supply (40)	0.1	14.64	29.3	61.0
Manufacturing of leather goods, luggage, and footwear (19)	0.1	10.84	43.1	39.2
Insurance and pension funding (66)	0.1	16.76	22.5	10.0
Manufacturing of coke and refined petroleum products (23)	0.0	17.81	48.7	65.1
Total contribution of the bottom 10 sectors	1.4			
Average across all sectors of the business economy		10.39	39.8	42.3

Table 1.9: Contribution of young firms to the aggregate real wage decline

This table quantifies the contribution of diverging wage trends at young firms relative to mature firms since 2002. The first three columns show the mean real wage, employment share, and percentage wage gap relative to mature firms for 4 firm age categories in the first half of the sample period (2002-2008): firm births, young firms 2-3 years old, young firms 4-9 years old, and mature firms aged 10 or more. ‘Not truly young’ firms are included in the mature category. The next three columns show the same variables for the second half of the sample period (2009-2015). The seventh column shows the percentage change in the mean real wage from the first to the second period for each firm age category. The last two columns illustrate the following thought experiment: how much of the aggregate decline in real wages between these two periods would have been avoided if the percentage wage gap between mature firms and the other firm age categories had remained constant?

Firm age	Actual							Counterfactual	
	Period 1 (2002-2008)			Period 2 (2009-2015)				Mean real wage	% change
	Mean real wage	Employment share	% Gap vs. 10+ firms	Mean real wage	Employment share	% Gap vs. 10+ firms	Real wage % change		
10+ years	15.67	76.9%	n.a.	15.38	79.4%	n.a.	−1.8%	15.38	−1.8%
4-9 years	15.80	14.4%	0.8%	13.46	13.2%	−12.5%	−14.8%	15.51	−1.8%
2-3 years	14.50	5.9%	−7.4%	12.45	5.1%	−19.1%	−14.2%	14.21	−2.0%
Firm births	13.75	2.8%	−12.2%	11.84	2.3%	−23.0%	−13.9%	13.46	−2.1%
Total	15.56	100.0%		14.89	100.0%		−4.3%	15.29	−1.7%
							Difference vs. actual		2.6pp
							% Improvement over actual		60%

1.9. Appendices

Appendix A: Sample construction

This appendix provides additional details on the construction of the sample and shows that the data used in this paper reflects several stylised facts of the overall UK business economy.

The sample is constructed by matching employee data from ASHE to employer data from the BSD. This is done based on a unique firm reference number provided by the Office for National Statistics (ONS). Table 1.A1 describes the construction of the sample. I focus on sectors C through K in the Standard Industrial Classification (SIC 2003) system, which cover mining, energy, manufacturing, construction, trade, finance, transportation, accommodation and food, real estate, renting, and business activities.¹ Throughout the paper, I call these sectors the ‘business economy’ for short, in line with the terminology used in ONS firm surveys. I exclude sectors like public administration, education, and health, since these are dominated by public employers, potentially leading to distortions in the wages paid in those sectors. The starting point in Column (1) is, therefore, the ASHE sample of 1.4 million annual observations on employees with a reported hourly wage working in the business economy. I then exclude employees on junior/trainee rates and those whose earnings in the reference period were affected by absence in Column (2), to avoid potential distortions in reported wages. This is a standard exclusion in all published aggregates derived from ASHE.

The resulting sample of 1.3 million job-year observations is matched to firm data to identify the observations pertaining to jobs at young firms. 98.9% of the firm reference numbers match to the same-year April BSD snapshot. The reasons for a less than 100% match are twofold. First, there is a small timing mismatch between when the snapshot of

¹I use the SIC 2003 classification, as it is supplied consistently in all years in the dataset.

Table 1.A1: Sample construction

This table contains observation counts by year for different stages of constructing the sample. The starting point (1) is the ASHE sample of employee jobs in the business economy with a reported hourly wage. Column (2) shows the sample count for employees whose pay was not affected by absence and did not correspond to a junior or trainee rate in the sample reference period. People whose earnings were affected in this way are excluded from the sample. This is also a standard exclusion in published aggregates based on ASHE data. Column (3) is the sample that matches to firm age data in the BSD in the same year. The match was done based on the unique firm reference number (entref variable), available in ONS microdata. Since a timing mismatch between the collection of employee and employer data is possible, column (4) adds the employee observations that match to next year's BSD snapshot of firms. The final column (5) shows the number of observations in the final sample. A breakdown of observations by firm age is shown at the bottom. 'Not truly young' firms are defined in Section 1.2.4.

Number of Observations					
Year	Start	Match to BSD in year t and t+1		Finish	
	(1)	(2)	(3)	(4)	(5)
2002	101,565	92,114	89,875	475	90,350
2003	100,324	91,564	90,565	429	90,994
2004	99,946	90,267	89,316	544	89,860
2005	103,056	95,816	94,794	526	95,320
2006	102,994	96,069	95,001	637	95,638
2007	89,655	83,902	83,088	416	83,504
2008	88,653	82,525	82,094	236	82,330
2009	102,979	96,272	95,788	229	96,017
2010	103,486	97,285	96,609	134	96,743
2011	109,571	102,885	101,996	366	102,362
2012	106,895	99,856	98,536	313	98,849
2013	107,903	100,708	99,561	594	100,155
2014	112,738	105,563	104,511	439	104,950
2015	110,933	103,545	102,441		102,441
Total	1,440,698	1,338,371	1,324,175	5,338	1,329,513
BSD-ASHE match %			98.9%	0.4%	99.3%
Observations at young firms:					
Aged 1 year			10,402	1,443	11,845
Aged 2 years			14,037	129	14,166
Aged 3 years			15,121	0	15,121
Total young firms			39,560	1,572	41,132
Aged 4-9 years			88,775	5	88,780
Aged 10+ years ('mature firms')			1,162,910	2,950	1,165,860
'Not truly young'			32,930	811	33,741

the firm data is taken (in April) and when ASHE data is received from firms (between May and June). Second, delays in administrative data feeds may lead a firm to appear in the BSD only some time after one its employees has been sampled for ASHE. This may be of

particular concern for newly-formed firms.² I therefore carry out a second match using the following year's BSD snapshot. The fourth column of Table 1.A1 shows that the second match adds only 5,338 observations, but 27% of them are at firms aged 1 year—a higher proportion than in the same-year match (where fewer than 1% of observations belong to firms of that age). The resulting sample has 1,329,513 job-year observations on 269,258 individuals working at 131,099 firms, for a total match rate of 99.3%. There are 28,344 unique individuals observed working at young firms, 243,760 individuals observed working at mature firms, and 17,131 individuals observed working at both types of firms during the sample period. Most of the observations in the final sample relate to mature firms, which can be expected given that the largest employers in the economy are also among the oldest.³

The sample reflects several stylised facts of the overall UK business economy. This is expected, given the random sampling of employees for ASHE and the fact that relatively few observations are lost in the matching process. Figure 1.A1 shows that the matched sample reflects a similar trend in the number of firms in the economy as the aggregate BSD dataset and published statistical time series. Consistent with this, Figure 1.A2 shows that the entry rate of firms in the matched sample also matches the aggregate patterns closely. Figure 1.A3 compares the sectoral breakdown of employment in the matched sample to that of the BSD. Services are somewhat under-represented in the sample, while industry and construction are over-represented. However, the trends in sectoral composition are

²Firms are legally obliged to register for VAT when they exceed an annual sales threshold (£82,000 or approx. US\$122,000 in 2015). Firms with sales below the threshold can register voluntarily. In addition, any firm with at least one employee paid more than £112 (US\$166) a week is legally required to register on the PAYE system. In 2015, the PAYE registration threshold equated to an adult on the minimum wage working 17 hours/week. A firm need only meet one of the VAT or PAYE conditions to be captured in the register. Therefore, only the very smallest firms are excluded. While exact estimates are not available, the ONS believes that the excluded businesses are largely self-employed individuals, rather than firms with non-owner employees (ONS 2006).

³A close match is not surprising: both datasets share national payroll tax data as one of their sources, and they include unique reference numbers specifically designed by the ONS for linking firm observations across datasets.

similar. The sample reflects the aggregate patterns of rising employment in services over this period, a decline in industrial employment, and a reversal in the pre-crisis growth trend of construction employment. Finally, Figure 1.A4 shows that the trends in the average real weekly earnings of the individuals in the sample are also a close reflection of the aggregate series. Earnings grew at a relatively constant, albeit slow, pace until the financial crisis. They then declined rapidly, as documented in Gregg et al. (2014), before stabilizing at a lower post-crisis level and rebounding slightly towards the end of the sample period.

Figure 1.A1: Firm count in the dataset vs. published aggregate time series

This figure compares the number of active firms in the population of firms in the business economy, captured in the BSD and in the matched sample, to the total firm count published in the ONS Business Demography. For better comparability of trends, firm counts are normalised to 100 in 2005, the first available year of the published series.




Figure 1.A2: Entry rate in the dataset vs. published aggregate time series

This figure compares the entry rates in the population of firms in the business economy captured in the BSD and in the matched sample to overall entry rates published in the ONS Business Demography. Each year's entry rate is calculated as the number of entrants in that year divided by the total number of active firms in the prior year. For better comparability of trends, the entry rates are normalised to 100 in 2006, the first year when the published aggregate was made available. In the case of the BSD and the matched sample, entrants are greenfield entrants, defined as described in Section 1.2.4.




Figure 1.A3: Sectoral composition of firms in the BSD and in the matched sample

This figure compares the high-level sectoral composition of the population of active firms in the business economy, captured in the Business Structure Database, to firms included in the matched sample.





Figure 1.A4: Average earnings in the matched sample vs. published aggregate time series

This figure compares the trend in average real private-sector earnings in Great Britain, published by the Office for National Statistics, to the year fixed effects from a regression of individual real hourly earnings on year dummies in the matched sample. Both series are deflated by the Consumer Price Index and indexed to 100 in 2002, the first year of the sample period. The evolution of mean real earnings in the sample tracks published data closely.



Comparing the matched sample to the BSD is also a convenient way to assess whether non-response rates in ASHE differ systematically for young and mature firms. Table 1.A2 shows the number of employees reported at firms captured in the BSD and the number of employees observed in the matched sample by firm size and age bands. While overall, ASHE appears to under-sample the employees of young firms (3.1% of the total vs. 5.1% in BSD), this is primarily because of the under-sampling of *small* firms. The potential impact of this on the results is mitigated by granular controls for firm size included in all the regressions with full controls, which results in comparisons of wages at young and mature firms being made within tight size bands.

Table 1.A2: Analysis of non-response in ASHE

This table shows the number of employees reported at firms in the BSD and the number of employees observed in the matched sample by firm size and age band. Data is pooled over 2006-2015, as the count of employees in the BSD is only available from that year onwards (prior to 2006 firms report employment, which includes working proprietors).

	Firm Size (Employment)				
Firm Age	< 10	10-99	100-249	250+	Total
<i>BSD</i>					
≤ 3 Years (Young)	2.7%	1.7%	0.3%	0.4%	5.1%
4-9 Years	3.9%	3.6%	0.8%	1.7%	10.0%
10+ Years (Mature)	6.4%	14.9%	6.2%	57.4%	84.9%
Total	13.0%	20.2%	7.3%	59.5%	100.0%
<i>Matched sample</i>					
≤ 3 Years (Young)	1.6%	1.0%	0.2%	0.3%	3.1%
4-9 Years	2.4%	3.1%	0.7%	1.3%	7.5%
10+ Years (Mature)	4.9%	16.9%	7.2%	60.5%	89.5%
Total	8.9%	21.0%	8.1%	62.1%	100.0%


Appendix B: Job creation and job destruction at young firms

This appendix contrasts the substantial contributions of young firms to job creation with their contributions to job destruction. It shows that young firms in Great Britain create disproportionately more gross jobs than their small size would suggest, relative to mature firms, and that they consistently create net jobs throughout the business cycle. These patterns are consistent with those observed in the US and other OECD countries (Haltiwanger et al. 2013; Criscuolo et al. 2014). At the same time, the failure rates of young firms are very high, leading to significant instability of the jobs that they create.

Contributions of young firms to job creation

Figure 1.B1: Contribution of young and mature firms to job creation and job destruction

This figure shows the shares of employment, job creation, and job destruction by firm size and age band. Within each firm size-age cell, job creation for each year from 2002 to 2015 is calculated as the sum of all positive year-over-year employment changes for firms that grew employment (including firm births). Similarly, job destruction is the sum of all negative year-over-year employment changes for firms that reduced employment or exited. Surviving firms are classified into size categories using on average employment, based on Haltiwanger et al. (2013). The shares are annual averages from 2002 to 2015.



Young firms create substantially more jobs relative to their size than mature firms. The bar chart in Figure 1.B1 illustrates this by dividing all firms in the BSD into 15 firm size-age

cells, and by plotting the share of total employment, job creation, and job destruction for each cell. These 15 cells are created, first, by dividing all private-sector firms in the BSD into three size bands: very small (employment less than 5), small (employment between 5 and 499), and large firms (employment of 500 or more). Second, within each size band, firms are classified into five age groups: births (aged 1 year or less), young (2 to 3 years old), young (4 to 9 years old), mature (aged 10 or more), and ‘not truly young.’ For each year of the sample period and each firm size-age cell, the number of jobs created is calculated as the total year-over-year increase in employment at firms that experienced net employment growth. Similarly, job destruction is the total year-over-year decrease in employment at shrinking firms. Each cell’s share of employment, job creation, and job destruction is then calculated for each year from 2002 to 2015. The bars represent average shares over this period.⁴

⁴The resulting figure parallels Figure 1 in Haltiwanger et al. (2013), who studied whether small firms or young firms contributed most to job creation in the US between 1992 and 2005. Unlike that paper, I separate ‘not truly young’ firms into a distinct category. Separating those firms from young firms strengthens the conclusions about the contributions of young firms to job creation.

I use broad firm size bands and age groups for clarity. Table 1.B1 demonstrates that similar conclusions about the substantial contribution of young firms to job creation hold when the data is divided into more granular size and age groups. It shows that ‘excess job creation’ (calculated for each firm age-size cell as that cell’s share of total job creation share minus its share of total employment) is positive at young firms and negative at mature firms, across detailed age and size categories.

Table 1.B1: ‘Excess’ contribution of young and mature firms to job creation

This table shows ‘excess’ job creation for detailed firm age-size cells, averaged across all years in the sample from 2002 to 2015, based on the BSD. Excess job creation in a given year is defined as the difference between that cell’s share of job creation and its share of total employment in that year (in percentage points). The analysis is based on 21,048,543 firm-year observations. Certain cells have been omitted due to disclosure restrictions. Positive numbers are concentrated among firms aged 1-5 years across most size categories, indicating that young firms contribute more to job creation than their share of total employment would suggest.

Firm Age	Firm Size (Employment)											
	1-4	5-9	10-19	20-49	50-99	100-249	250-499	500-999	1,000-2,499	2,500-4,999	5,000-9,999	10,000+
1 (Firm Births)	11.5	2.0	1.4	1.4	0.6	0.5	0.3					
2	2.4	0.5	0.4	0.4	0.3	0.2	0.1	0.2				
3	0.5	0.4	0.3	0.3	0.2	0.2	0.1	0.1				
4	-0.1	0.2	0.2	0.2	0.1	0.1	0.1	0.1				
5	-0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.1				
6-10	-0.7	0.0	0.0	0.2	0.2	0.3	0.2	0.1	0.1			
11-15	-0.7	-0.3	-0.3	-0.2	-0.1	-0.0	0.0	0.1	0.1	0.1	-0.0	-0.1
16-20	-0.5	-0.4	-0.4	-0.4	-0.3	-0.2	-0.1	-0.1	0.0	-0.1	-0.1	-0.2
21-25	-0.4	-0.3	-0.4	-0.5	-0.3	-0.3	-0.2	-0.2	-0.1	-0.2	-0.2	-0.4
26+	-0.7	-0.8	-1.1	-1.6	-1.2	-1.4	-1.2	-1.1	-1.7	-1.6	-1.4	-9.5

Figure 1.B1 leads to two main results about the role of young firms as an engine of job creation. First, firm births and firms that are 2 to 3 years old account for a greater share of total job creation relative to their share of employment. The difference between these two shares, which can be called the ‘excess’ contribution to job creation (Haltiwanger et al. 2013), is observed across all size bands. Firms aged 4 to 9 years also contribute disproportionately to job creation, except in the smallest size band, which captures firms with employment of less than 5. In contrast, the share of jobs created by mature firms is smaller than these firms’ share of employment in each size category. Second, young firms aged 2 to 3 and 4 to 9 years account for a greater share of job creation than of job destruction. (This figure is based on annual data, so firm births cannot destroy any jobs, as they had no employment in the prior period.) In contrast, mature firms’ share of jobs destroyed is substantially greater than their share of jobs created. Finally, firms classified as ‘not truly young’ behave very similarly to mature firms. The disproportionate contribution of ‘not truly young’ firms to job destruction suggests that they may be appearing as new entities in the firm register as a result of a restructuring of existing businesses, rather than as a result of being greenfield start-ups. The pattern of job destruction supports the decision to classify these firms as a category distinct from young firms in the remainder of this paper.

Naturally, the absolute number of jobs created each year is higher at mature firms than at young firms. This is because mature firms tend to be larger, and therefore, in aggregate, account for a high share of the aggregate employment stock.⁵ Nevertheless, the data underlying Figure 1.B1 highlights the disproportionate contribution of young firms to job creation: on average, between 2002 and 2015, firms up to 3 years old accounted for 29.2% of total annual job creation, despite accounting for only 7.4% of total employment.

⁵There are only 91 firm births and 420 firms 2 to 3 years old in the “large” category, suggesting that few young firms reach the threshold of 500 employees rapidly.

Do young firms create more jobs than they destroy? The bar chart with job creation and job destruction shares in Figure 1.B1 is not best-suited to answer this question. This is because if total employment in the economy is shrinking in a given year, then a firm size-age cell might, in fact, be losing jobs on a net basis, even if its share of total job creation exceeds its share of total job destruction.

To answer the question whether young firms create more jobs than they destroy, I calculate ‘net job creation’ as the difference between the number of jobs created and the number of jobs destroyed for each year between 2002 and 2015, and for five groups: firms aged 2 to 3 years, 4 to 5 years, 6 to 9 years, 10 or more years, and ‘not truly young’ firms. In other words, I take the net within-firm job creation and job destruction measures and net them further across firms within each age group. I exclude firm births, because by definition, they do not destroy any jobs relative to the prior period. In addition, I break down the 4 to 9 age group into two smaller groups: firms that are 4 to 5 years old and firms that are 6 to 9 years old. The purpose of this is to show more clearly how net job creation varies with firm age. The annual net job creation figures calculated on this basis are plotted in Figure 1.B2.

Young firms have consistently created jobs on a net basis, even during the Great Recession. Across all age groups, net job creation is highly cyclical. However, it is consistently positive at young firms and particularly high at firms aged 2 to 3. On the other hand, mature firms consistently destroy net jobs, and ‘not truly young’ firms behave similarly to mature firms through the cycle. Table 1.B2 repeats the analysis of net job creation for more granular firm size and age cells, with similar conclusions.

Figure 1.B2: Net job creation by young and mature firms through the cycle

This figure shows net job creation by firm age and year from 2002 to 2015, based on the BSD. Net job creation is defined as jobs created minus jobs destroyed, as described in the main text. The left-hand chart plots firms aged 2-3 years, 4-5 years, and 6-10 years. The right-hand chart plots firms aged more than 10 and ‘not truly young’ firms. The vertical axis shows net job creation in thousands. While net job creation is cyclical, young firms have consistently created net jobs over this period, even during the recession following the 2008 financial crisis. In contrast, mature firms have consistently destroyed jobs over the entire sample period.




Table 1.B2: ‘Excess’ contribution of young and mature firms to *net* job creation

This table shows ‘excess’ *net* job creation for detailed firm age-size cells, averaged across all years in the sample from 2002 to 2015, based on the BSD. Excess net job creation in a given year is defined as the difference between that cell’s share of net job creation (job creation minus job destruction) and its share of total employment in that year (in percentage points). The analysis is based on 21,048,543 firm-year observations. Certain cells have been omitted due to disclosure restrictions. Positive numbers are generally concentrated among firms aged 1-10 years across most size categories with only a few exceptions, indicating that young firms contribute more to net job creation than their share of total employment would suggest.

Firm Age	Firm Size (Employment)											
	1-4	5-9	10-19	20-49	50-99	100-249	250-499	500-999	1,000-2,499	2,500-4,999	5,000-9,999	10,000+
1 (Firm Births)	206.2	37.6	24.4	23.4	7.9	7.0	2.8					
2	49.9	5.6	1.5	0.2	-0.4	1.4	3.2	3.4				
3	20.0	8.8	7.0	6.5	3.5	5.0	2.3	0.2				
4	9.9	7.1	4.9	4.7	2.4	2.6	1.7	0.4				
5	5.2	5.4	3.8	3.8	0.7	1.5	1.5	2.2				
6-10	9.6	12.0	10.8	12.3	5.4	6.7	5.1	3.5	6.0			
11-15	-27.7	-3.8	-2.6	-2.4	-2.4	-3.0	-2.2	2.6	-2.3	1.8	0.3	-5.2
16-20	-17.2	-4.9	-4.8	-4.5	-3.0	-2.6	-1.5	0.4	3.9	4.8	6.3	1.4
21-25	-11.1	-3.8	-4.4	-4.3	-3.4	-1.4	-1.2	-4.6	2.5	-7.2	1.8	-5.6
26+	-20.3	-10.2	-12.2	-13.8	-9.6	-12.0	-13.3	-4.8	-1.4	4.6	10.1	-15.6

These analyses confirm the importance of young firms to job creation highlighted in the recent literature. My results on excess job creation are consistent with the conclusions of Haltiwanger et al. (2013) for the US between 1992 and 2005, as well as with cross-country results reported in Criscuolo et al. (2014). By focusing on a new British dataset, this paper contributes to confirming that the capacity of young firms to create a disproportionate amount of jobs relative to their size can be observed in different countries and time periods. Moreover, my conclusions on the positive contribution of young firms to net job creation during a period that included the Great Recession suggest that the ability of young firms to create more jobs than they destroy in aggregate are resilient over time. The prominent role that young firms play in aggregate job creation highlights the importance of assessing what types of jobs these firms create.

The risk of joining a young firm as an employee


One question of particular relevance for the employees who join young firms is how long those jobs can be expected to exist. Are the jobs at young firms more or less risky than jobs at mature firms? While the analysis in the previous section shows that young firms, as a group, continue to create net jobs after birth, their contribution to net job creation each year is the difference between a large number of jobs added and a large (albeit smaller) number of jobs destroyed. For instance, the data underlying the plot of net jobs created for firms aged 2-3 years shows that these firms added 3.8 million jobs between 2002 and 2015, but also destroyed 1.5 million jobs during the same period (based on within-firm netting of annual job creation and destruction).

One of the main ways in which young firms destroy jobs is through exits. In this section, ‘exit’ denotes a complete shutting down of a firm’s operations, rather than, for instance, a sale to another firm. One motivation for disregarding other forms of exit is that it is not clear

when the sale of a firm denotes success or failure in the data. In general, for many young firms, particularly for those that aim to grow and innovate, the sale to a larger firm is a sign of success for the entrepreneur and the initial investors in the firm. However, the owners of a failing firm may also choose to sell as an alternative to shutting down, often quickly and for a low price. The data does not allow me to distinguish between a successful sale and a fire sale. In both cases, some of the original firm's establishments may continue operating, and its employees may continue working for the new owners. Although I can deduce that a sale has occurred, I do not observe the transaction price or any other details about the sale process, other than the mere fact of an ownership change. I therefore focus on firm closures, given the risk that they pose to the job stability of the employees of young firms.⁶

Figure 1.B3: Proportion of jobs destroyed by firm exits

This figure shows the cumulative 1-year, 2-year, and 3-year exit rates for firms aged 1 year (firm births) and firms aged 10 or more in the BSD. Exits are defined as firm liquidations, as discussed in the main text. The exit rates are weighted by employment to reflect the proportion of the initial stock of jobs destroyed by firm exits. Data is based on 7,423,126 firm-year observations from 2002 to 2012. It is not possible to calculate a 3-year exit rate for subsequent years, as the sample ends in 2015.



Young firms have high exit rates compared to mature firms, and almost 40% of the jobs

⁶The closure of a firm and its establishments may not necessarily represent a failure of the business plan, either. It could simply be driven by the owner's decision to retire or pursue another activity that generates a higher return. However, if none of the establishments continue operating, then it is likely that the firm's employees will need to search for another job. Hence, firm closure is a relevant source of risk for those individuals.

created by a young firm in its first year of life are destroyed by an exit within the next 3 years. In contrast, mature firms that destroy jobs do so mainly by downsizing. Figure 1.B3 illustrates this by plotting the cumulative 1-year, 2-year, and 3-year employment-weighted exit rates of 1-year-old young firms next to the cumulative 1-year, 2-year, and 3-year exit rates of mature firms aged 10 or more. For each year from 2002 to 2012, the employment-weighted exit rate for each firm age category is calculated by dividing (a) total employment in year t of firms that will have exited by year $t + 1$ by (b) total year t employment of all firms in that age category. The bars in Figure 1.B3 represent annual averages of the exit rates over the 2002-2012 period. Weighting by employment means that these exit rates can be interpreted as the proportion of jobs destroyed due to firm exit within the next 1, 2, and 3 years. The figure shows that 15.6% of jobs created by young firms in the first year of their life are destroyed through firm exit by the second year, and a cumulative 38.3% are destroyed by the fourth year. Mature firms, on the other hand, only destroy 4.0% of their initial stock of jobs through exit within the next 3 years.⁷


The fact that mature firms have lower exit rates than young firms does not necessarily mean that jobs at mature firms are more stable. For instance, if mature firms make systematically worse matches with workers, then the average employee who joins a mature firm might even face a similar probability of separation than one who joins a young firm. To see whether this is the case, I use the data to estimate whether it is more or less likely that an employee of a young firm leaves her job, compared to an employee of a mature firm.

To determine whether a separation from a job is more or less likely at a young firm than at

⁷For reference, Figure 1.B4 in shows firm exit rates not weighted by employment for the same firm age categories and period as Figure 1.B3. The exit rates of firms aged 1 are similar to those in Figure 1.B3. This is because exiting young firms are not systematically larger or smaller than non-exiting young firms. On the other hand, the unweighted firm exit rate of mature firms aged 10 years or more is substantially higher than the employment-weighted exit rate of firms in this age category. This is because small mature firms exit at a much higher rate than large mature firms. The employment-weighted rates are more relevant, as they illustrate the proportion of jobs lost. Nevertheless, Figure 1.B4 shows that even on an unweighted basis, the exit rate of young firms is substantially higher than the exit rate of mature firms.

Figure 1.B4: Unweighted firm exit rates

This figure shows the cumulative 1-year, 2-year, and 3-year exit rates for firms aged 1 year (firm births) and firms aged 10 or more, based on the BSD. Exits are defined as firm liquidations, as discussed in the main text. Unlike in Figure 1.B3, the exit rates are not weighted by employment. Data is based on 7,423,126 firm-year observations from 2002 to 2012. It is not possible to calculate a 3-year exit rate for subsequent years, as the sample ends in 2015.



a mature firm, I run probit regressions of the probability of exiting a job on firm age dummies with the full of controls relating to firm, job, and individual characteristics available in the dataset (the controls are described in Section 1.3.2 of the main text). The results are shown in Table 1.B3. The table considers all separations in the sample, independently of the reason (e.g., firing or a voluntary departure), which is not available in the data. In Column (1), the dependent variable is a binary indicator of whether the employee leaves the current job by the following period. In Column (2), the dependent variable is a binary indicator which equals 1 when the employee leaves the current job by the following period, and he or she is observed as an employee at another firm. The objective of this is to check whether the results change when transitions out of the dataset (to unemployment, self-employment, retirement, etc.) are excluded. The table shows the relative probability of leaving a job for an employee at a firm aged 3 or less, compared to a firm aged 10 or more.

The probit regressions suggest that job instability is indeed greater at young firms than at mature firms. The probability of separation is 12.3% higher at a young firm, and 7.9% higher

Table 1.B3: Job instability at young firms

This table reports the marginal effect of working for a firm aged 3 years or less on separation in the next period, relative to working for a firm aged 10 years or more. The results are based on probit regressions that include ‘full controls,’ as described in Section 1.3.2 in the main text. Standard errors are clustered at the firm level.

	Separation	Separation and observed as an employee next year
	(1)	(2)
Firm aged ≤ 3 years	0.123*** (0.004)	0.079*** (0.004)
Number of Observations	1,099,036	947,082

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

when conditioning on the departing employee joining another firm next period, compared to a mature firm. These estimates are highly significant. Therefore, even though the results cannot be interpreted causally, they show that jobs at young firms are statistically more likely to lead to separations than jobs at mature firms. This is consistent with the results of Brixy et al. (2011) for Germany and Goetz, Hyatt, McEntarfer, and Sandusky (2017) for the US. The balance of evidence suggests that joining a young firm as an employee is a risky proposition.

Appendix C: Robustness checks

This appendix discusses robustness checks pertaining to the results in Sections 1.3 to 1.5 of the main text. It begins by considering the relative wages of new hires at young and mature firms using various measures of compensation and various wages to define the sample of young firms and their employees. It then tests the sensitivity of the results on wage returns. In all, this appendix shows that the results discussed in the main text hold for a range of alternative sample and variable definitions.

Starting wages

Table 1.C1 compares the compensation of new hires at young and mature firms using two alternative measurements. Panel A defines compensation as total hourly compensation, which adds up wages and employer pension contributions calculated on a per-hour basis. If young employers contribute less to their employees' pension pots than mature employers, then this may affect the estimate of the wage gap between the two categories of firms. Indeed, Figure 1.C1 shows that firms aged less than 10 are less likely to provide pension benefits than firms aged 10 or more.⁸ Panel B defined compensation as weekly earnings. To the extent that hours worked differ systematically between firms of different ages, the results based on weekly earnings could differ from those based on hourly wages. Focusing on column (7), which reports results for the full sample of employees with individual fixed effects, both total hourly compensation and weekly earnings are approximately 1% higher for new hires at young firms relative to mature firms. This is consistent with the results for hourly wages in the main text. Panel C shows that when comparing similar jobs and individuals, hours worked do not differ substantially between young and mature firms.

⁸Figure 1.C1 highlights a number of interesting trends in the provision of defined-benefit and defined-contribution pension benefits in Great Britain. These trends are discussed in Adrjan and Bell (2018).

Table 1.C1: Robustness to alternative measures of employee compensation

This table reports the coefficient on the dummy for a firm aged 3 years or less from regressions of log weekly earnings (Panel A) and total weekly hours (Panel B). Column headings correspond to the scenarios described in the notes to Tables 1.3 and 1.4. Standard errors clustered at the firm level are in parentheses.

	Raw data	New hires	Full controls	Displaced workers	Individual selection on unobservables		
					Labor market entrants	Individual fixed effects (labor market entrants)	Individual fixed effects (all workers)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Total Hourly Compensation							
Firm aged ≤ 3 years	-0.156*** (0.016)	0.022* (0.013)	0.004 (0.004)	0.026 (0.021)	0.010* (0.006)	0.025*** (0.007)	0.018*** (0.003)
R ²	0.03	0.02	0.71	0.74	0.53	0.62	0.43
Number of Observations	1,329,513	272,969	179,812	2,208	56,089	165,428	1,226,228
Panel B: Weekly Earnings							
Firm aged ≤ 3 years	-0.272*** (0.025)	-0.015 (0.026)	-0.009 (0.005)	0.015 (0.030)	0.008 (0.012)	0.007 (0.012)	0.011*** (0.004)
R ²	0.01	0.01	0.74	0.78	0.69	0.59	0.42
Number of Observations	1,329,513	272,969	179,812	2,208	56,089	165,428	1,226,228
Panel C: Weekly Hours							
Firm aged ≤ 3 years	-2.774*** (0.301)	-0.713** (0.362)	-0.064 (0.081)	-2.587 (2.139)	-0.054 (0.162)	-0.364** (0.181)	-0.077 (0.060)
R ²	0.01	0.01	0.69	0.44	0.76	0.44	0.44
Number of Observations	1,329,513	272,969	179,812	2,208	56,089	165,428	1,226,228

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.C2: Robustness to alternative definitions of the sample

This table reports the estimated new hire premium at young firms relative to mature firms from regressions of log hourly wages and log weekly earnings after imposing a series of restrictions on the sample to test the robustness of the results. All regressions include full controls, as described in Section 1.3.2. Standard errors clustered at the firm level are in parentheses.

	Domestic- owned firms	Excl potential owners	Excl postcodes with 100+ firms	Excl small firms	Excl tech and finance
	(1)	(2)	(3)	(4)	(5)
Panel A: Labor Market Entrants					
Hourly Wages	0.014** (0.006)	0.008 (0.006)	0.012** (0.006)	0.010* (0.006)	0.012** (0.006)
Weekly Earnings	0.014 (0.012)	0.004 (0.013)	0.010 (0.012)	0.009 (0.012)	0.007 (0.012)
Number of Observations	44,838	55,169	55,789	55,755	51,716
Panel B: Individual Fixed Effects (Labor Market Entrants)					
Hourly Wages	0.026*** (0.007)	0.015* (0.008)	0.021*** (0.007)	0.020*** (0.007)	0.022*** (0.007)
Weekly Earnings	0.012 (0.013)	0.003 (0.014)	0.009 (0.012)	0.008 (0.012)	0.005 (0.013)
Number of Observations	126,031	162,689	164,613	164,922	141,495
Panel C: Individual Fixed Effects (All Workers)					
Hourly Wages	0.019*** (0.003)	0.009*** (0.003)	0.014*** (0.003)	0.020*** (0.003)	0.017*** (0.003)
Weekly Earnings	0.020*** (0.004)	0.005 (0.005)	0.011*** (0.004)	0.017*** (0.004)	0.016*** (0.004)
Number of Observations	944,487	1,183,035	1,221,689	1,220,106	1,021,841

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1.C1: Proportion of jobs with pension benefits

This figure plots the proportion of jobs with pension benefits for firms aged less than 10 and firms aged 10 or more, distinguishing between defined-benefit pensions (left panel) and defined-contribution pensions (right panel). Grey solid lines denote the upper and lower bounds of 95% confidence intervals.




Table 1.C2 shows the sensitivity of the estimated wage gap between new hires at young and mature firms to different ways of defining young firms and their employees. The table is divided into three panels. Panel A shows the results of robustness checks for the subsample of labour market entrants, which corresponds to the specification in Column (2) of Table 1.4. Next, Panel B shows the results for labour market entrants with individual fixed effects, which corresponds to Column (3) of Table 1.4. Finally, Panel C shows the results with individual fixed effects for all workers, which corresponds to Column (4) of Table 1.4.

The conclusion that young firms pay a 1-2% premium to new hires, compared to mature firms, is robust to a number of alternative ways to define young firms and their employees. Column (1) excludes foreign-owned firms from the sample to remove any potential confounding effects of young firms that are subsidiaries of multinational enterprises. Such firms may transfer existing managers from other countries or implement established compensation policies that differ from those at home-grown start-ups unconnected to existing

corporate groups. Column (2) deals with the potentially serious concern that some of the employees in the dataset could be firm owners, rather than employees. While the ASHE questionnaire specifically asks firms to exclude firm owners from the data, it is nevertheless possible that some owners who receive a salary through the PAYE income tax system are present in the dataset. To address this, I test the robustness of the results to dropping the first observed employee hired in the first year of a firm's life. Column (3) excludes postcodes with 100 or more registered firms, which may suggest tax-driven firm creation by self-employed individuals.⁹ In a related vein, Column (4) excludes any firms that only ever achieve maximum employment of 1 over their lifetime, on the basis that some of these may represent self-employed individuals rather than employees working for someone else. Column (5) considers the issue that wages and earnings in ASHE do not capture equity compensation. To the extent that young firms are more likely to use equity compensation, the estimates of the initial wage premium at those firms may be underestimated. One way to approach this is to remove firms in the technology and finance sectors, where equity compensation is more common than in other sectors. While precise data on the prevalence of equity vs. wage compensation by sector is not available for the UK, US data shows that tech is by far the leader in the extent of equity compensation, measured in various ways (Fredrick W. Cook & Co., Inc. 2015). The results for the remaining (non-tech, non-finance) sectors, are shown in Column (5).¹⁰

Looking across all five columns, the results do not change significantly relative to the

⁹A typical UK postcode is much smaller than a US zip code, with an average area of less than 1 km². In urban areas, a postcode often corresponds to a single building.

¹⁰I use the Eurostat definition of the tech sector. It comprises "high-tech manufacturing" sectors chosen based on technological intensity, defined as R&D expenditure/value added (24.4-Manufacture of pharmaceuticals and medicinal chemicals; 30-Manufacture of office machinery and computers; 32-Manufacture of radio, television and communication equipment; 33-Manufacture of medical, precision and optical instruments, watches and clocks; and 35.3-Manufacture of aircraft and spacecraft), and "knowledge-intensive services" based on the share of tertiary-educated individuals by 2-digit sector (64-Post and telecommunications; 72-Computer and related activities; and 73-Research and development), according to the SIC 2003 classification. I exclude sector 64.11-National post activities.

main sample discussed in the main text. Focusing on Panel C, the conclusion remains that wages and earnings at young firms are 1-2% higher than at mature firms. These results are still statistically significant.

Table 1.C3: Accounting for firm fixed effects

This table reports the difference in the log hourly wage and log weekly earnings between the employees of firms aged 3 years or less and the employees of firms aged 10 years or more, based on regressions that include both individual and firm fixed effects. All models include full controls from Column (5) of Table 1.3, as described in the main text. Standard errors clustered at the firm level are in parentheses.

	Labor market entrants	All individuals
	(1)	(2)
Panel A: Hourly Wages		
Firm aged ≤ 3 years	0.028** (0.015)	0.023*** (0.004)
R ²	0.89	0.93
Number of Observations	149,643	1,180,092
Panel B: Weekly Earnings		
Firm aged ≤ 3 years	0.035 (0.023)	0.033*** (0.006)
R ²	0.90	0.93
Number of Observations	149,643	1,180,092

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$


Next, Table 1.C3 checks whether the conclusions are affected when unobserved, time-invariant firm heterogeneity is controlled for through firm fixed effects. The results are shown for labour market entrants in Column (1) and the full sample of employees in Column (2). The specification corresponds to Column (4) in Table 1.4, with the addition of firm fixed effects, and the coefficients represent the difference in the wages of new hires at young firms relative to mature firms. The results suggest a wage premium at young firms of 2.8% for labour market entrants and 2.3% for all individuals. This is higher than the estimates without firm fixed effects, suggesting that firm heterogeneity does play a role, but the difference is not statistically significant. The results are similar to the estimate of a 2.5% premium obtained when comparing new hire wages at young firms that attain the top

quartile of productivity conditional on survival to similarly (observably) productive mature firms.

Figure 1.C2: Returns to joining a young firm based on weekly earnings

This figure presents the evolution of weekly earnings of individuals hired by a firm aged 3 years or less, compared to individuals hired by a firm aged 10+ years. Panel A shows the subsequent evolution of earnings regardless of whether the individuals stay at the same firm or are observed at other firms (unconditional returns). Panel B shows returns to tenure within continuing matches. It plots the coefficients on tenure dummies from a regression of individual weekly earnings on tenure dummies interacted with firm age dummies, full controls, and individual fixed effects. In each plot, earnings are expressed relative to initial new hire wages at mature firms, represented by the horizontal baseline. The thin grey lines represent 95% confidence intervals. There is no evidence that the employees hired by young firms subsequently experience higher average earnings growth, compared to the employees hired by mature firms.

Panel A: Unconditional Returns



Panel B: Returns in Continuing Matches






Figure 1.C3: Robustness of estimated returns to working at a young firm (continuing matches)

This figure repeats the analysis of returns to tenure in a continuing job match from Figure 1.5, excluding potential entrepreneurs and young firms that had a foreign owner in the first observed year. A potential entrepreneur is defined as the first individual observed working at a given firm and hired in the first year of the firm's life.

Panel A: Hourly Wages



Panel B: Weekly Earnings



Wage growth

Figure 1.C2 plots the unconditional returns to joining a young and a mature firm, as well as the returns in continuing matches, using weekly earnings as an alternative measure of compensation. The results are similar to those using the hourly wage.

Figure 1.C3 applies selected further robustness checks to the graphical analysis of the estimated returns to working at a young firm within a continuing match. Of the sensitivities discussed in the context of Table 1.C2, excluding potential firm owners from the sample had the largest absolute effect on the estimate of the young firm wage premium, while dropping foreign-owned firms led to the largest reduction in the sample size. Figure 1.C3 shows that applying both of these restrictions does not change the conclusion about wage growth at young and mature firms, although the standard errors are now substantially wider.


Wage evolution over time

Similarly, Figure 1.C4 checks the sensitivity of the evolution of the estimated young firm premium or penalty over time to several key restrictions. One of the key features of the data discussed in the main text is that relative wages at young firms appear to have fallen rapidly since 2012, compared to mature firms. To check whether this is related to the increase in self-employment in the UK over the same period, Panel A excludes firms with maximum observed employment of 1, while Panel B drops potential firm owners, and Panel C excludes postcodes with 100 or more firms. All these restrictions aim to address potential issues with self-employed individuals registering as employees of their own firms. The plots show that the general pattern of falling relative wages at young firms is unaffected.

Overall, the key results discussed in the paper are robust to a range of restrictions on the sample and alternative definitions of key variables.

Figure 1.C4: Robustness of the estimated young firm wage premium/penalty by year

This figure shows the average wage premium at firms aged 3 years or less relative to firms aged 10 years or more. The vertical axis represents the marginal effect of working for a young firm by year for a new hire, based on individual regressions of log hourly wages on firm age dummies interacted with year dummies. All regressions include “full controls,” as described in the main text, and individual fixed effects. These charts show that the negative trend in wages at young firms relative to wages at mature firms, which remains after controlling for observables and individual selection on time-invariant unobservables, cannot be explained by (a) the smallest firms, (b) the dataset potentially capturing firm owners rather than employees, or (c) postcodes with a large number of firms, which may suggest outsourcing or tax-driven firm creation.



Appendix D: Additional figures and tables

Figure 1.D1: Distribution of hourly wages and log hourly wages by firm age

This figure plots the distribution of real hourly wages and log real hourly wages by firm age category, considering all jobs and pooling across all years in the sample from 2002 to 2015. Wages are deflated by the Consumer Price Index and expressed in 2014 pounds sterling (GBP). Panel (a) shows the density of the real hourly wage, right-censored at £25 for clarity. The total number of observations is 1,181,951. Panel (b) shows the density of the log real hourly wage, right-censored at 4 (approximately £57). The total number of observations is 1,308,644. This figure shows that the wage distribution at firms aged between 4 and 9 years lies between the wage distributions at firms aged 3 years or less and firms aged 10 years or more. As a result, the discussion in the main text focuses mainly on the differences between the two extreme firm age categories.




Figure 1.D2: Distribution of hourly wages by firm age in 2015

Panel (a) compares the relative distribution of log hourly wages at firms aged 3 years or less with wages at firms aged 10 years or more in 2015, the last year of the sample. Wages at young firms are more concentrated at and around the adult minimum wage, which applied to individuals aged 21 or more. Wages at mature firms are relatively more concentrated in the middle and at the higher end of the wage distribution. The distribution is right-censored at £25/hour for clarity, resulting in 3,326 observations at firms aged 3 years or less and 82,735 observations at firms aged 10 years or more. Panel (b) shows data only for individuals aged 21 or more, to demonstrate that observations below the adult minimum wage in panel (a) relate almost exclusively to younger individuals, who face a lower (or no) statutory minimum wage. There are 2,964 observations at firms aged 3 years or less and 76,159 observations at firms aged 10 years or more.





Figure 1.D3: Distribution of weekly earnings and hours by firm age

This figure compares the relative distribution of real weekly earnings and hours at firms aged 3 years or less (shaded bars) with real weekly earnings and hours at firms aged 10 years or more (transparent bars). Data is pooled over all years in the sample (2002-2015). Earnings are deflated by CPI and expressed in 2014 pounds. Panel A shows all jobs, while panel B shows new hires only.

Panel A: All Employees



Panel B: New Hires Only





Figure 1.D4: Correlation between wages and firm size

This figure illustrates the relationship between the average log real hourly wage and firm size, where firm size is measured by total employment. Panel (a) shows the estimated firm size coefficients from a regression of individual real hourly wages on 99 centiles of employment, where the omitted category is the first centile (firms with employment of 1). Each point represents between 7,853 and 20,288 person-year observations at firms in a given centile. Panel (b) shows the same estimated firm size coefficients after controlling additionally for the 5-digit sector. Vertical bars indicate 95% confidence intervals based on clustering by firm. The total number of observations is 1,328,987.

(a) Raw correlation between average wages and firm size



(b) Estimated firm size fixed effects after controlling for the 5-digit sector




Figure 1.D5: Age distribution of individuals with and without a prior observed wage

This figure shows the employee age distribution for observations in the matched sample with a prior observed wage (solid line) and observations without a prior observed wage (dashed line).




Table 1.D1: Moments of the wage distribution

This table shows the characteristics of the real hourly wage and the log real hourly wage distributions by firm age category in the matched sample, pooled across all years from 2002 to 2015. Panel A shows that within each firm age category the mean wage is substantially higher than the median. Dispersion is high and rises with firm age. In addition, since the distribution becomes more right-skewed as firms age, the distance between the mean and the median also increases. Panel B, on the other hand, shows that the distribution of log wages is less dispersed and less skewed. Moreover, dispersion and skewness do not vary substantially with firm age. As a result, the distance between the mean and the median is constant across firm age categories.

(a) Real hourly wage

	Firm age category			
	≤ 3 years	4-9 years	10+ years	Not true young
Mean	12.94	14.18	14.92	15.81
Median	9.69	10.69	11.25	12.06
Standard deviation	11.20	12.63	14.69	13.35
Skewness	8.60	8.53	58.97	7.65
Mean - median	3.25	3.50	3.67	3.75
Number of observations	41,132	88,780	1,165,860	33,741

(b) Log real hourly wage

	Firm age category			
	≤ 3 years	4-9 years	10+ years	Not true young
Mean	2.38	2.47	2.52	2.57
Median	2.27	2.37	2.42	2.49
Standard deviation	0.53	0.55	0.55	0.57
Skewness	1.08	1.03	0.97	0.78
Mean - median	0.11	0.10	0.10	0.08
Number of observations	41,132	88,780	1,165,860	33,741

Table 1.D2: Alternative methods of controlling for firm size

This table uses the specification of equation (1.1) with full controls and individual fixed effects in Column (4) of Table 1.4 to explore different methods of controlling for firm size. For comparison purposes, Column (1) controls for 19 ventiles of employment, like in Table 1.4. Column (2) controls for 99 centiles of employment. Column (3) adopts a linear specification typical in the literature on wages at young firms. Similarly to Burton et al. (2017), I find that controlling for size in such an inflexible way changes the conclusions about the effect of firm age on wages. Column (4) uses a third-degree polynomial. Standard errors clustered by firm are in parentheses.

	Dependent variable: log(real hourly wage)			
	Ventiles (1)	Centiles (2)	Linear (3)	Polynomial (4)
Firm aged ≤ 3 years	0.014*** (0.003)	0.017*** (0.003)	0.003 (0.003)	0.015*** (0.003)
Firm aged 4-9 years	0.011*** (0.002)	0.011*** (0.002)	0.009*** (0.002)	0.011*** (0.002)
Not truly young	0.006 (0.005)	0.006 (0.004)	0.006 (0.005)	0.007 (0.005)
log(1 + employment)			0.012*** (0.000)	0.073*** (0.005)
log(1 + employment) ²				-0.009*** (0.001)
log(1 + employment) ³				0.000*** (0.000)
F-statistic (all firm size variables)	86	32	715	393
Observations	1,226,228	1,226,228	1,226,228	1,226,228

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.D3: Productivity distribution by firm age

This table shows the relative distribution of observations in the matched sample by quartile of the overall distribution of real sales per employee, a proxy for labour productivity. Young firms are more concentrated in the bottom quartile of the distribution. Nevertheless, they are also represented in the top quartile.

	Quartile of the overall productivity distribution				Total	Observations
	Bottom 1	2	3	Top 4		
Firms aged ≤ 3 years	55.9	25.9	10.5	7.7	100.0	41,132
Firms aged 4-9 years	42.0	26.9	16.2	14.9	100.0	88,780
Firms aged 10+ years	22.5	24.8	26.4	26.4	100.0	1,165,860
'Not truly young' firms	31.9	24.9	19.8	23.5	100.0	33,741

Table 1.D4: Wages of new hires as a predictor of a young firm's success

This table reports the results of individual-level regressions, where the dependent variable is a measure of firm success that is either: a) a 0/1 indicator of firm survival (columns 1 and 2) or b) the log of sales per worker ("productivity," columns 3 and 4). The independent variables are measured in year t , while the dependent variables are measured in year $t + x$, where $x = 1$ or 3, as indicated above each column. All regressions include individual fixed effects and "full controls," as described in the main text. Standard errors are clustered at the firm level.

	Measure of firm success (dependent variable)			
	Survival (0/1)		Log(Productivity conditional on survival)	
	1-year horizon (1)	3-year horizon (2)	1-year horizon (3)	3-year horizon (4)
(a) Log(Real hourly wage)	−0.000 (0.001)	−0.001 (0.003)	0.193*** (0.011)	0.173*** (0.010)
(b) Firm aged ≤ 3 years	−0.008 (0.007)	−0.038** (0.018)	0.043 (0.061)	0.115** (0.062)
(c) Log(Real hourly wage) * Firm aged ≤ 3 years	0.006** (0.003)	0.017** (0.007)	−0.196*** (0.027)	−0.092*** (0.025)
(a+c) Marginal effect of higher wages on success for young firms	0.006** (0.003)	0.017** (0.007)	−0.003 (0.026)	0.081*** (0.025)
Observations	1,167,661	965,404	1,159,479	910,391

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Bibliography

- Brixy, Udo, Susanne Kohaut, and Claus Schnabel. 2011. "Employment Stability in Newly Founded Firms: A Matching Approach Using Linked Employer-Employee Data from Germany." *Small Business Economics* 36(1): 85-100.
- Criscuolo, Chiara, Peter Gal, and Carlo Menon. 2014. "The Dynamics of Employment Growth: New Evidence from 18 Countries." OECD Science, Technology and Industry Policy Papers No. 14. Paris: OECD Publishing.
- Frederic W. Cook & Co., Inc. 2015. *Aggregate Share-Based Compensation Report. Aggregate Share-Based Compensation Across Industries and Size*. https://www.fwcook.com/content/Documents/Publications/The_2015_Aggregate_Share-Based_Compensation_Report.pdf (last accessed 10 May 2018).
- Goetz, Christopher, Henry Hyatt, Erika McEntarfer, and Kristin Sandusky. 2017. "The Promise and Potential of Linked Employer-Employee Data for Entrepreneurship Research," in J. Haltiwanger, E. Hurst, J. Miranda, and A. Schoar (eds.), *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, Chicago: University of Chicago Press.
- Haltiwanger, John, Ron Jarmin, and Javier Miranda. 2013. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95 (2): 347-361.
- Office for National Statistics. 2006. *Business Structure Database User Guide*. <https://discover.ukdataservice.ac.uk/catalogue?sn=6697> (last accessed 16 March 2017).

Chapter 2

The Mightier, the Stingier: Firms' Market Power, Capital Intensity, and the Labour Share of Income*

Abstract

What determines the proportion of a firm's income that workers receive as compensation? This paper uses longitudinal firm data from a period of substantial labour share variation to understand the firm-level determinants of the labour share of income—a question that has typically been addressed with country- and sector-level data. Firms with greater market power and a higher ratio of capital to labour allocate a smaller proportion of their value added to workers. These results are consistent with a simple model of the determinants of the labour share, in which profit-maximising firms employ labour and capital to produce a final good in an imperfectly competitive market. They suggest that firm-level drivers play a key role in the evolution of the aggregate labour share, which has declined significantly since the 1970s.

Keywords: Labour Share, Employee Compensation, Market Power, Capital Intensity

JEL Codes: D33, E25, J24, J30

* I thank Brian Bell, Steve Bond, Doug Gollin, Ines Lee, Ferran Mañé (discussant), Matteo Richiardi, Luis Valenzuela Rivera, and seminar participants at the Royal Economic Society PhD Meetings in London, 6th PhD Student Workshop on Industrial and Public Economics in Reus, and the Royal Economic Society Annual Conference in Brighton for helpful comments and suggestions.

2.1. Introduction

The classic economic question of how rents are distributed among the factors of production has recently returned to the fore, motivated by a substantial decline in labour's share of national income in many of the world's economies since the 1970s (Rodríguez and Jayadev, 2010; Elsby, Hobijn, and Şahin, 2013; Karabarbounis and Neiman, 2014; IMF, 2017). Figure 2.1 illustrates this decline for the G7 countries. From a starting point of approximately 70% in 1970, the aggregate labour share of income in this group of advanced economies has declined by an average 2 percentage points per decade. This persistent downward trend contrasts with the historical stylized fact of long-run stability of the labour share, which was prominently noted by Keynes (1939) and Kaldor (1957) and has since become a standard assumption in many micro- and macroeconomic models.

Several potential causes of the decline in the global labour share have been proposed. One leading account is that technological change or declines in the price of capital relative to labour have led firms to substitute capital for labour, thus decreasing the overall share of income that accrues to the latter factor of production (Acemoglu, 2003; Bentolila and Saint-Paul, 2003; Karabarbounis and Neiman, 2014). Another is that deregulation or other drivers have increased the market power of firms, raising the profit share of income at the expense of the labour share (Blanchard and Giavazzi, 2003; Azmat, Manning, and Van Reenen, 2012; Barkai, 2016; Autor, Dorn, Katz, Patterson, and Van Reenen, 2017; De Loecker and Eeckhout, 2017). To date, however, both explanations have only been tested at the country and sector level, even though the aggregate labour share is a result of the production decisions and wage-setting processes occurring at individual, heterogeneous firms. In this paper, I propose a methodology for calculating the labour share of value added at the firm level and use UK data to show that the labour share varies significantly across firms, even within narrowly-defined sectors, as well as within firms

over time. I then investigate what determines the labour share at individual firms. I find that firms with greater market power and a higher ratio of capital to labour allocate a smaller proportion of their value added to workers. These results contribute to a better understanding of the drivers of changes in aggregate factor income shares.

To motivate why the labour share might vary across and within firms, I describe a simple model of a firm that employs capital and labour to produce a final good in an imperfectly competitive product market. When the optimal combination of inputs is chosen to maximise profits, the labour share is determined by the firm's market power and the capital intensity of production. Market power matters for the labour share because if a firm can set the price of its product at a higher mark-up over cost, then a greater share of value added will go to profits, at the expense of labour. Capital intensity matters too, because if a production process is highly automated and requires little labour input, then under certain conditions, a higher proportion of the firm's value added will be allocated towards compensating the suppliers of capital. In the model, the exact nature of the relationship between capital intensity and the labour share depends on the elasticity of substitution between labour and capital and the nature of technical progress. Whether the labour share is higher or lower at firms that employ more capital relative to labour is therefore an empirical question.

To test whether a firm's market power and capital intensity impacts the labour share, I use a large longitudinal dataset of firms in the United Kingdom covering the period from 2005 to 2012, during which the aggregate labour share experienced substantial fluctuations. Thanks to mandatory filing requirements that affect nearly all incorporated businesses in the UK, it is possible to calculate the labour share of value added for a wide universe of firms, including small and medium enterprises, operating across a variety of sectors. Using the model to motivate some potential drivers of this heterogeneity, I find an inverse

relationship between a firm's labour share and its market power (measured by the market share), and between the labour share and the capital intensity of production (measured by the capital-labour ratio). To account for potential endogeneity that may result, for instance, from a joint response of variables to unobserved shocks, I estimate dynamic specifications of the model using the Generalised Method of Moments (GMM). The finding that the labour share is inversely related to the firm's market power and capital intensity is robust to alternative econometric methods, definitions of key variables, sample selection decisions, and empirical specifications.

The results of this paper have several policy implications. As technological developments and capital accumulation drive many economies to become more capital intensive, my estimates suggest that labour's share of aggregate income may continue to decline. Similarly, the labour share might remain under pressure if economic activity becomes increasingly concentrated in a smaller number of firms that wield greater market power at the expense of the consumer, as some authors have suggested is the case (Barkai, 2016; Autor et al., 2017; De Loecker and Eeckhout, 2017). Moreover, a declining labour share implies a rising disconnect between average productivity and pay within the economy. This is because the labour share of value added can be expressed as the ratio of average compensation to the average product of labour. The link between productivity and pay matters greatly to households, because wages and salaries represent by far the largest source of household income in most developed countries. Understanding the determinants of firm-level labour shares is therefore important for assessing the potential impact of future productivity gains or losses on the earnings of the average household.

Does the finding that higher capital intensity of production lowers the labour share of income mean that in the aggregate, capital and labour are substitutes? This result is consistent with some strands of the factor share literature. For instance, Karabarbounis and

Neiman (2014) estimate an elasticity of substitution between capital and labour greater than one in their cross-country analysis. In addition, Bentolila and Saint-Paul (2003) find a negative relationship between the labour share and the capital-to-output ratio in sector-level data from selected OECD countries. Piketty's (2014) explanation of the global evolution of aggregate income shares of capital and labour also relies on the high substitutability of these two factors of production. On the other hand, an inverse relationship between capital intensity and the labour share appears to be at odds with a large part of the literature that focuses specifically on estimating the elasticity of substitution between capital and labour. The value of this elasticity continues to be a matter of debate. While available estimates vary widely depending on the methodology and context, often exceeding one, many authors find lower values that imply a far lesser degree of substitutability (see Chirinko, 2008, for a survey, and Knoblach, Rößler, and Zwerschke, 2016, for a meta-analysis of US studies).

One possible reason for the divergent estimates of the aggregate elasticity of substitution is that different classes of capital and labour are substitutable with one another to a different extent. To understand whether this might drive my overall results, I carry out two tests. First, I divide the sample into high-wage and low-wage firms, based on average compensation per employee. I find that the negative effect of capital intensity on the labour share is driven mainly by low-wage firms. Under the assumption that the average wage level at a firm is a proxy for the average skill level of the workforce, these results imply that it is mainly low-skilled labour that is substitutable for capital, rather than high-skilled labour. Second, I consider the effects of tangible and intangible capital separately, based on the book value of long-term tangible and intangible assets reported in the financial statements of firms. Tangible assets are physical items such as machinery, buildings, and land, while intangible assets represent nonphysical concepts such as patents, trademarks, and goodwill. Accounting standards place limitations on the types of intangible assets that

can be recognised on the balance sheet (particularly those that are internally-generated; see Haskel and Westlake, 2017, for a discussion), and such assets are often difficult to value. Nevertheless, I find evidence that an increase in intangible assets per employee has a positive effect on the labour share of high-wage firms. This suggests low substitutability between intangible capital and high-skilled labour. The outcomes of both tests are consistent with intuition provided by the hypothesis of capital-skill complementarity (Griliches, 1969). The paper thus contributes to bridging the gap between the two opposing views on the potential impact of increases in capital intensity on the labour share.

My results complement the macroeconomic literature on factor shares by focusing on the determinants of the labour share at individual firms. Existing theoretical models of aggregate labour share determination ascribe a leading role to factors that, in practice, vary across narrowly-defined sectors and individual firms, such as the characteristics of production technologies, relative quantities of capital and labour inputs, and the extent of monopoly power in the product market (Bentolila and Saint-Paul, 2003; Azmat et al., 2012; Barkai, 2016; Autor et al., 2017; De Loecker and Eeckhout, 2017). The paper's contributions to this literature are to: a) construct firm-level measures of these variables using panel microdata, b) show that the labour share varies substantially across firms, as well as within firms over time, and c) exploit this variation to test hypotheses previously proposed in the context of a representative firm and measured only using aggregate data. A firm-level focus is motivated by the fact that in developed economies, most economic activity outside the public sector is formally organized in firms. In the UK, for instance, the corporate sector represents approximately 65% of aggregate gross value added, while private-sector employees account for 67% of total employment (ONS, 2016). The sharing of income between labour and other factors of production thus occurs primarily at this level.

Understanding how the labour share at individual firms is determined is therefore helpful for understanding the distribution of aggregate national income.

This paper also extends a small body of recent research that has begun to use firm-level data to analyse aggregate labour share trends. Growiec (2012) uses data on Polish firms to study the impact of entry/exit behaviour, ownership, and market conditions on labour share fluctuations. Böckerman and Maliranta (2012) use plant-level data from Finland to decompose changes in the aggregate labour share into within-plant changes on one hand, and changes in the industry structure due to entry, exit, and reallocation of activity across firms on the other hand. The focus of these two papers is on aggregate labour share dynamics, whereas I assess the role of firm and market characteristics in determining the level of the labour share at individual firms. Siegenthaler and Stucki (2015) exploit a survey of 4,000 Swiss manufacturing, construction and business service firms to estimate correlations between the labour share and a firm's use of computers, number of competitors, modern organizational arrangements, participation in export markets, bargaining coverage, and the share of female employees. In contrast to that paper, I use a large panel of firms operating in all sectors of the business economy to test two leading hypothesized drivers of labour share determination, capital intensity and market power. Moreover, I use methods designed to address the potential endogeneity or simultaneity biases that may arise in estimating the magnitude and direction of these relationships. Finally, two recent papers have focused on the potential relationship between market power and the labour share. De Loecker and Eeckhout (2017) document a rise in average mark-ups of US firms and provide suggestive graphical evidence of a correlation of this trend with the decline in the aggregate US labour share. Autor et al. (2017) use firm data from several countries to construct measures of industry concentration and relate them to sector-level labour share trends. The authors build a model, in which highly productive “superstar”

firms capture an increasing market share, leading to a fall in the aggregate labour share through a reallocation of activity to such firms. Imposing a Cobb-Douglas production function, which precludes any role for capital intensity, they document an empirical association between rising industry concentration and declining labour shares within sectors. My paper complements the analysis of these two papers in two main ways. First, it assesses the relative effect of various determinants of the labour share at the level of individual firms, rather than examining sectoral trends. Second, it allows both capital intensity and market power to affect a firm's labour share, allowing the magnitude of these two effects to be compared. The role of capital intensity is particularly relevant, given the threat of labour-to-capital substitution in many industrialized markets and its potential to put further downward pressure on the aggregate labour share in the future. Overall, my results suggest that to understand the underlying causes of the trends in the aggregate labour share, we should look closely to firms, and the environment that they operate in, for answers.

The remainder of the paper is structured as follows. Section 2.2 discusses how to measure the labour share of income at the firm level and documents the dispersion of labour shares across firms and within firms over time. A simple model of the determinants of the firm-level labour share is discussed in Section 2.3. Section 2.4 describes the data, and Section 2.5 outlines the empirical strategy. Section 2.6 presents the results and tests their robustness, while Section 2.7 discusses the role of human capital and the firm's labour market power in determining the labour share. Section 2.8 concludes.

2.2. Measuring the labour share of income at the firm level

In this section, I discuss two types of contextual information relevant for the analysis of the determinants of the labour share. First, I discuss how to measure the labour share of income using data from firms' financial reports in a way that mirrors the underlying economic concept. Then, I show that firm-level labour shares are highly dispersed, even within narrowly-defined sectors, which is inconsistent with the predictions of a fully competitive model. This observation motivates the remainder of the paper, which explores the role of market power and capital intensity in labour share setting.

The labour share is defined as the ratio of labour cost to some measure of income. In standard formulations based on national accounts, the numerator typically includes both wage and non-wage compensation of employees. In the denominator, income is usually measured by aggregate gross value added.¹

$$\text{Labour Share} = \frac{\text{Labour Compensation}}{\text{Income}} = \frac{\text{Wage} \times \text{Labour}}{\text{Price of Value Added} \times \text{Real Value Added}}$$

A well-documented problem with this approach is that it does not account correctly for self-employed individuals. While the output of the self-employed is part of national income, the labour component of their remuneration is not captured in the compensation of employees. This understates the labour share and affects cross-country comparisons (Gollin, 2002).²

In this paper, I focus on the firm-level labour share, using data on incorporated businesses. One advantage over macroeconomic data is that focusing on firms and their

¹ In the SNA 2008 national accounting framework, compensation of employees includes wages and salaries, paid in cash or in kind, as well as employer contributions to pensions, healthcare, and social insurance. Gross value added is often used as a measure of income when the labour share is calculated for individual sectors or regions. This is because GDP includes taxes on products, such as VAT, which are only known for the whole economy. In the UK, GVA accounts for approximately 90% of GDP (ONS, 2016).

² Guerriero (2012) reviews a range of adjustments to the aggregate labour share proposed in the macroeconomic literature. The income of sole proprietors and unincorporated businesses ("mixed income") represents approximately 7% of total GVA in the UK (France 7%, Germany 10%, United States 10%).

employees removes the confounding effects of self-employment on labour share estimates. Although some self-employed individuals may in fact operate as incorporated businesses, their impact on the data is mitigated by showing that the conclusions of this paper hold when the smallest firms are excluded.

At the firm level, I define the labour share as the ratio of total employment expense to gross value added. Like in the national accounts, employee compensation reported by individual firms includes non-wage expenses, such as pension and health insurance expenses, as well as social security taxes. I estimate firm-level gross value added as sales minus the cost of intermediate inputs other than employee compensation (calculated from the financial statements as earnings before interest, tax, depreciation, and amortization plus employee compensation expense), adjusted for inventory growth:

$$\begin{aligned}\text{Firm-Level Labour Share} &= \frac{\text{Employee Compensation}}{\text{Gross Value Added}} \\ &= \frac{\text{Employee Compensation}}{\text{Sales} - \text{Intermediate Inputs} + \Delta \text{Inventories}}\end{aligned}$$

Adding the change in inventories to the denominator aligns the firm-level measure of value added with the corresponding economic concept of revenue-based output. Consider what would happen if inventory changes were ignored. A firm that has produced goods in a given period but has not yet sold them will have accounted for the full cost of labour, capital, and other inputs used in production. Sales net of the cost of intermediate inputs will thus be low, and the labour share will be overstated. Conversely, reported sales can reflect the running down of inventories of final goods produced in earlier periods.

The adjustment for inventory changes is consistent with the way aggregate gross value added is calculated in the national accounts (ONS, 2016). The existing literature on

firm-level labour shares has not incorporated this adjustment into its measures of gross value added. However, year-to-year inventory changes at individual firms can be large. Excluding them from the calculation of gross value added can affect the conclusions about the aggregate evolution of the labour share. I demonstrate this in Section 2.4, in the context of the data sample.

When the numerator and denominator are divided by the number of employees, the labour share becomes a ratio of the average wage to the average product of labour:³

$$\text{Firm-Level Labour Share} = \frac{\frac{\text{Employee Compensation}}{\text{Number of Employees}}}{\frac{\text{Gross Value Added}}{\text{Number of Employees}}} = \frac{\text{Average Compensation}}{\text{Average Product of Labour}}$$

The lower the labour share, the greater the gap between wages and average productivity. Firms that have low labour shares thus undercompensate workers relative to the amount of output they create. A declining labour share thus implies a “decoupling” of wages and productivity and a shift in how the benefits of economic activity are shared with workers.

What does basic economic theory predict for the variation of labour shares across firms? In a framework of competitive markets and free mobility of homogeneous input factors, all profit-maximizing firms with identical, well-behaved production functions will choose the same technologically-efficient input allocation. In equilibrium, each factor will be paid its marginal product, and the labour shares will be equal across firms. Nothing will be gained by using firm-, rather than sector-level, data to analyse the determinants of the labour share.

³ Average wages and productivity are expressed on a per-employee basis, as I do not have firm-level data on hours worked. Pessoa and Van Reenen (2013) analyse the relationship between aggregate wage and productivity growth in the UK using a variety of measures and macroeconomic data sources.

However, the data shows wide dispersion in firm labour shares. Table 2.1 presents measures of between-firm and within-firm variation in the labour share and its components, both in levels and in logs, derived from the random-effects model, $x_{it} = \alpha + u_i + \varepsilon_{it}$ applied to the pooled sample. Between-firm variation ($\sigma_u/|\bar{x}|$) is measured as the standard deviation of the error term u_i normalized by its mean, while within-firm variation ($\sigma_\varepsilon/|\bar{x}|$) is measured by the standard deviation of ε_{it} , normalized by its mean. When the labour share is measured in levels, between-firm variation is 0.40, and within-firm variation is 0.60, such that 31% of the variance is due to differences across firms. The corresponding measures are 1.14 and 1.08 when the labour share is measured in logs, with 53% of the variance arising between firms. The empirical strategy in this paper exploits both sources of variation.

The simplest framework that predicts no dispersion of the labour share also predicts no dispersion in its individual components. Table 2.1 shows that both average compensation and the average product of labour vary widely in the sample, and most of this variation is across firms in the same industry. The fact that the ratio of these two variables is also dispersed means that the numerator and denominator do not correlate perfectly. One way to see this is to regress average compensation and the average product of labour and calculate $1 - R^2$ to obtain the proportion of total variation in the former that is not explained by the latter. Controlling for year and 4-digit sector fixed effects, this proportion is 54% when the variables are measured in levels and 42% when the variables are measured in logs. The remainder of this paper investigates the sources of this variation.

2.3. Why might capital intensity and market power matter for the labour share?

This section outlines a simple model to illustrate why the labour share may differ across firms within a given sector, even under the assumption that the production technologies are identical. The aim is to motivate the empirical analysis that follows in an intuitive manner before proceeding to a discussion of the data and the empirical strategy.⁴

To see why a firm's labour share might be related to its market power and the capital intensity of production, consider a profit-maximizing firm i with the production function $Y_i = Y(A_i L_i, B_i K_i)$, where Y_i is output, L_i is labour, K_i is capital, and A_i and B_i are labour- and capital-augmenting productivity, respectively, which can vary across firms. Labour and capital are supplied elastically at the wage w and rental rate r , but the product market is imperfectly competitive.⁵ From the first-order condition for labour, the labour share is

$$s_i^L \equiv \frac{wL_i}{P_i Y_i} = \varepsilon_i^{YL} \left(1 + \frac{1}{\eta_i} \right) \quad (2.1)$$

where $\varepsilon_i^{YL} = \frac{\partial Y_i}{\partial L_i} \frac{L_i}{Y_i}$ is the partial elasticity of output with respect to labour and $\eta_i = \frac{\partial Y_i}{\partial P_i} \frac{P_i}{Y_i}$ is the price elasticity of demand. The first term, ε_i^{YL} , is a function of the productivity parameters and the labour and capital inputs. The firm's labour share thus depends on the relative productivity-augmented quantities of the factors employed in production. The sign of the relationship between the labour share and the relative quantities of capital and labour inputs is determined by the elasticity of substitution between capital and labour.⁶

⁴ For a discussion of similar predictions generated by general equilibrium models of monopolistic competition, see Karabarbounis and Neiman (2014) and Barkai (2016).

⁵ Appendix D contains supporting derivations of key equations in the paper. The assumption of a perfectly competitive labour market is relaxed in Section 2.7.

⁶ This point was first raised by Hicks (1932) and Robinson (1933). If the elasticity of substitution between capital and labour is exactly one (like in the model of Autor et al. (2017), who assume a Cobb-Douglas production function), then the partial elasticity of output with respect to labour, ε_i^{YL} , will be independent of the firm's capital-to-labour ratio. In that scenario, the empirical analysis that follows should find that capital per worker does not matter for a firm's labour share (which is not the case). Arpaia, Pérez, and Pichelmann

If returns to scale in production are constant, then the first term of equation (2.1) can be expressed more transparently as a function of the labour-to-capital ratio and the firm-specific productivity parameters:

$$\varepsilon_i^{YL} \equiv \frac{Y_L(A_i L_i, B_i K_i)}{Y(A_i L_i, B_i K_i)} L_i = \left(\frac{A_i L_i}{B_i K_i} \right) \frac{Y_L \left(\frac{A_i L_i}{B_i K_i}, 1 \right)}{Y \left(\frac{A_i L_i}{B_i K_i}, 1 \right)} = g \left(\frac{A_i L_i}{B_i K_i} \right) \quad (2.2)$$

where Y_L is the partial derivative of output with respect to labour. The empirical analysis that follows will therefore test whether there is a relationship between the relative factor inputs (measured as net tangible assets per employee) and the firm's labour share, with the ratio of productivity parameters absorbed by firm fixed effects.

The second term of equation (2.1) denotes the firm's product market power and is determined by the characteristics of demand. In the baseline specification of my empirical analysis, I use the firm's share of 4-digit sector sales to proxy for its product market power. This is motivated by the observation that some models of imperfect competition predict a direct relationship between the price elasticity of demand and firms' market shares in equilibrium. For example, consider a Cournot oligopoly with n firms, each producing a homogenous product with marginal cost c_i and facing a linear inverse demand curve, $P = a - bQ$, where $Q = \sum_i q_i$ is total quantity produced. In equilibrium, firm i 's market share is $s_i^M \equiv \frac{q_i}{\sum_i q_i} = \frac{n+1}{n} \frac{a-c_i}{a-\bar{c}} - 1$, where $\bar{c}$ is the average marginal cost. The market share can be expressed in terms of the elasticity of demand, η , as

$$s_i^M = \frac{(n+1)(a-c_i)}{a+n\bar{c}} (-\eta) - 1 \quad (2.3)$$

(2009) develop a model of the aggregate labour share where low- and high-skilled labour differ in their elasticity of substitution with respect to capital.

This illustrates a possible link between firms' market shares and the price elasticity of demand.⁷ The intuition is that when there are fewer firms in a sector, each firm's market share is higher. Lower competition leads to greater monopoly power and a lower total quantity produced, which corresponds to a more elastic part of the demand curve. This simple model motivates the empirical analysis that follows, which uses longitudinal firm data to estimate the firm-level determinants of the labour share.

2.4. Data

The main source of data for my analysis is the Financial Analysis Made Easy (FAME) database supplied by Bureau van Dijk. It contains the financial filings of UK firms that submit annual reports to Companies House, a government agency responsible for maintaining the official company register. Virtually all companies in the UK, except the smallest ones, are legally required to file an audited annual report. In 2015, the compliance rate with this requirement was 99.1% (Companies House, 2015). Financial information follows standardized accounting conventions. The audit requirement helps ensure that the data is of high quality, partially mitigating potential concerns about reporting error. For these reasons, FAME data is routinely used in research (e.g., Draca, Machin, and Van Reenen, 2011).

Thanks to strict reporting requirements, FAME contains information on a large number of private and unlisted companies. Such companies represent the bulk of the total

⁷ In the empirical model, firm fixed effects absorb any firm-specific cost advantage, which will affect the relationship between market share and equilibrium price elasticity of demand in this model if marginal costs are heterogeneous and the differences are persistent. Note that if demand is iso-elastic, there will be no relationship between the (constant) price elasticity of demand and market shares as industry concentration changes. However, I show in Section 2.6 that the conclusions about the role of market power in determining the firm-level labour share hold in a number of different specifications.

firm count in the UK (and in most other countries), as well as a high proportion of aggregate employment. The inclusion of small, private firms represents both an opportunity and a challenge. On the one hand, it generates a more heterogeneous and representative sample of firms than databases of listed companies, which exclude most small and medium enterprises. On the other hand, many small firms are likely to be family businesses or self-employed individuals operating as incorporated businesses. The labour share of such firms can easily be mismeasured, since the labour and capital components of compensation are difficult to disentangle in the case of business owners. Section 2.6 considers this issue explicitly and shows that the results of this paper are robust to excluding the smallest firms from the sample.

The sample for analysis is defined by including all market sector firms that report information on labour compensation, sales, tangible assets, employment, and the line items needed to calculate value added. This requirement makes the firms in the sample larger, on average, in terms of sales and employment than the universe of UK firms. The financial sector is excluded due to methodological differences in the calculation of value added. I also exclude sectors with a substantial presence of non-market services (public administration, health, education, arts, etc.), where value added can be difficult to measure. To avoid capturing foreign subsidiaries of UK firms and their employees, I use unconsolidated financial accounts.

The result is a panel of 119,764 observations on 31,402 firms, covering the period from 2005 to 2012. This period was characterized by substantial variation in the aggregate labour share. Figure 2.2 shows that the aggregate labour share in the sample closely tracks the market sector labour share calculated from the national accounts, and follows a similar downward trend as the labour share in the UK economy and other G7 countries. (Notably, while the Great Recession led the labour share to peak in 2009 as the value of output

dropped precipitously, it does not seem to have fundamentally altered the general downward trend over this period.) The firms in the sample represent a substantial proportion of economic activity generated by the UK non-financial market sector. In the last year of the sample period, 2012, they generated total value added of £275bn, employed 5.8 million people, and spent £167bn on employee compensation – 29%, 27%, and 32% of the corresponding aggregate market sector figures reported by ONS (2018).

Table 2.2 reports basic descriptive statistics for firm sales, employment, and the variables used in the empirical analysis. In the table and in the main specification of the empirical model the labour share is calculated as total compensation expense divided by gross value added, as described in Section 2.2. Capital intensity is calculated as the book value of tangible fixed assets net of accumulated depreciation (‘net tangible assets’), divided by total employment. Tangible fixed assets are long-term assets such as machinery, buildings, and land. The market share is calculated as each firm’s share of total sales in the 4-digit sector. Nominal values are converted to constant prices using the GDP deflator and ratios are winsorised at the 1st and 99th percentile. The median firm reports sales of £12 million and employment of 73. It has a labour share of value added of 75%, £11,000 of net tangible assets per employee, and a market share of less than 1%.⁸ The sample includes firms of all sizes – from those employing a handful of individuals to the largest businesses in the economy with thousands of employees.

As mentioned in the discussion on measurement above, the year-to-year change in inventories can be very large for an individual firm, substantially affecting firm-level estimates of gross value added (and therefore the labour share). Inventory changes in the

⁸ I show later that the conclusions of this paper hold when using alternative measures of the labour share, capital intensity, and market power. I do not use the perpetual inventory method to calculate a measure of capital due to the short time series of data available for many of the firms in the dataset.

sample range from a minimum of -£1.4bn to a maximum of +£1.8bn. To see the importance of including inventory changes for the correct measurement of gross value added, consider the impact of the financial crisis, which represented the largest shock to aggregate output during the sample period. With the correct adjustment for inventory changes, aggregate nominal GVA in the sample fell by 5.4% from 2008 to 2009. This falls squarely between the 4.3% decrease reported in the national accounts for “private non-financial corporations” (ONS, 2016) and the 6.8% decrease reported in the aggregated results of the Annual Business Survey of firms in the “non-financial business sector” (ONS, 2018).⁹ On the other hand, excluding the inventory adjustment would have led incorrectly to an estimated nominal GVA *increase* in the sample of 0.3%, contrary to the observed trend. This shows that correct measurement of the labour share is important for drawing macroeconomic conclusions from firm-level data.

Table 2.2 also shows that, on average, a firm’s labour share declines as its capital intensity rises. Firms with tangible assets per employee of less than £10,000 have a median labour share of 81%. As the capital intensity rises across the sample, the average labour share drops precipitously. For instance, firms with £50,000 to £100,000 of tangible assets per worker have a median labour share of 66%, while firms with tangible assets per worker of £100,000 to £500,000 have a median labour share of 52%. In the right tail of firms with tangible assets per worker of £500,000 or more, the median labour share is only 20%. A similar association holds within broad sectors (manufacturing and services). While the median manufacturer is more capital intensive than the median service firm (£15,068 vs.

⁹ The Annual Business Survey is one of the sources used for calculating aggregate GVA in the national accounts, but the national accounts incorporate a number of other adjustments. See the discussion in Section 2.2 and ONS (2016) for more details.

£9,532 of tangible assets per worker, respectively), the labour share within each sector declines as capital intensity rises.

Similarly, the labour share declines with a firm's market share. Firms with market share within their 4-digit sector of less than 1% have a median labour share of 76%. The average labour share declines gradually for firms observed with higher market shares. Like in the case of capital intensity, the right tail of firms with market share of 20% or greater has the lowest median labour share of 65%. Figure 2.3 illustrates this relationship visually. Both capital intensity and market power thus appear to be potentially relevant determinants of the firm-level labour share.

2.5. Empirical strategy

My starting point for assessing whether the negative correlation between a firm's labour share and its capital intensity on one hand, and between the labour share and market share on the other, hold when controlling for other confounding variables and addressing potential sources of endogeneity is the following panel data model, based on log-linearized equation (2.1):

$$s_{it}^L = \beta_1 \left(\frac{K_{it}}{L_{it}} \right) + \beta_2 s_{it}^M + \theta_j + \varphi_r + \delta_t + (a_i + v_{it}) \quad (2.4)$$

where s_{it}^L is the labour share at firm i at time t ; $\left(\frac{K_{it}}{L_{it}} \right)$ is the capital-to-labour ratio and s_{it}^M is the market share. Technology is permitted to vary across sectors j . The firm-specific labour share is also allowed to be affected by factors specific to the region r where the firm is located, and by changes in the average labour share over time, denoted by the year fixed effects δ_t .

Two problems may arise in estimating equation (2.4) using ordinary least squares. One is the potential correlation between the error term and the covariates. For instance, both the capital-to-labour ratio and the labour share could be jointly determined in response to idiosyncratic shocks, v_{it} . An example of this might be a rise in the bargaining power of workers that leads to a higher labour share while also causing firms to substitute away from labour towards capital, inducing a positive correlation between capital intensity and the labour share. This will lead to biased results.

A second problem is that time-invariant managerial or organizational characteristics, captured by a_i , may be correlated with one or more regressors, such as the firm's choice of factor inputs or market position. For instance, there can be unobserved differences in how value added is shared among labour and other factors of production, which may depend on unobserved preferences of management, working conditions or other social and organizational factors, or the relative size of the firm-specific productivity parameters A_i and B_i . Insofar as those unobserved differences are permanent, they are captured by the time-invariant parameter a_i in equation (2.4). Reverse causality is also a concern. For instance, firms that have a lower labour share than peers ($a_i < 0$) may find it easier and cheaper to fund increases in their capital stock, thanks to their higher profits and the ability to offer a higher return on capital. A low labour share might thus cause higher capital intensity, not the other way around.

If the shocks v_{it} are serially uncorrelated, then the standard method for dealing with these problems involves estimating equation (2.4) in differences to eliminate the impact of permanent heterogeneity, a_i , and using instrumental variables to address concerns about the correlation between the regressors and the error term. While external instruments for firms' input choices are typically difficult to find, lagged values of the endogenous

regressors themselves, dated $t - 2$ or earlier, would be potentially valid instruments, as they would be uncorrelated with the error term.

However, in practice, when the model is estimated using this procedure, the residuals exhibit serial correlation beyond the first-order negative serial correlation that would be induced by first-differencing if the shocks were serially uncorrelated. This casts doubt on the validity of a static model and suggests a need to model serial correlation in the errors explicitly. But if the shocks are serially correlated, then the error term is a function of all past disturbances. In this situation lagged values of the endogenous regressors are no longer valid instruments. A model with firm fixed effects suffers from the same problem. Hence there is no consistent way to estimate the static model.

A standard approach to dealing with this problem in the GMM literature (e.g., Bond and Guceri, 2017) is to estimate a dynamic, autoregressive distributed lag model that contains the lagged labour share and lags of the regressors on the right-hand side and uses first-differencing to remove the time-invariant firm effect (see Appendix D for derivations):

$$\begin{aligned} \Delta s_{it}^L = & \rho \Delta s_{i,t-1}^L + \pi_1 \Delta \left(\frac{K_{it}}{L_{it}} \right) + \pi_2 \Delta \left(\frac{K_{i,t-1}}{L_{i,t-1}} \right) + \pi_3 \Delta s_{it}^M + \pi_4 \Delta s_{i,t-1}^M \\ & + \Delta \varepsilon_{it} \end{aligned} \quad (2.5)$$

In this model, if the unobserved shocks v_{it} are serially correlated and follow an AR(1) process, for instance, then lags of any endogenous regressor x_{it} dated $t - 2$ or earlier are now potentially valid instruments for x_{it} and $x_{i,t-1}$, as $E[x_{i,t-s}\varepsilon_{it}] = 0$ for $s \geq 2$. Similarly, the lagged dependent variable $s_{i,t-1}^L$ can be instrumented with lags of s_{it}^L dated $t - 2$ or earlier. The long-run relationship between capital intensity and the labour share is then given by $\frac{\pi_1 + \pi_2}{1 - \rho}$. Similarly, the long-run effect of market share on the labour share is

$\frac{\pi_3 + \pi_4}{1 - \rho}$. These parameters can be estimated consistently using two-stage least squares or the Generalised Method of Moments (GMM).

An alternative way to motivate a dynamic model with lags of the labour share, capital intensity, and market share is costly or slow adjustment of employment, investment, wages or sales to shocks, especially when adjustments are costly. In the data, firms' labour shares, capital-to-labour ratios, and market shares are highly persistent over time.¹⁰ This observation is consistent with either the presence of adjustment lags or serially correlated productivity shocks.

To estimate the coefficient vector $(\rho, \pi)'$, I use the System GMM estimator developed by Arellano and Bover (1995) and Blundell and Bond (1998). In addition to using equation (2.5) in first differences with lagged levels of the potentially endogenous variables as instruments, System GMM also exploits the corresponding equation (2.4) in levels, with lagged first-differences as instruments. This approach has two main advantages compared to applying GMM only to the equation in first differences. First, it uses more information by exploiting between-firm variation, which is an important source of overall variability in labour shares, besides within-firm variation. In the pooled sample, 53% of variation in the log labour share is between—rather than within—firms. Second, when the data are persistent, System GMM exhibits a smaller finite-sample bias (Blundell and Bond, 1998).

System GMM is the main approach used in the remainder of this paper. To avoid concerns with overfitting (Roodman 2009), I limit the number of instruments to the lagged levels dated $t - 2$ and $t - 3$ in the first-difference equation and lagged differences dated

¹⁰ Appendix A provides evidence on the persistence of these variables.

$t - 1$ and $t - 2$ in the levels equation, rather than using all past values.¹¹ The key additional assumption is that the first-differences of the explanatory variables (as well as the first-differences of the labour share) are uncorrelated with the time-invariant unobserved heterogeneity captured by α_i and not transformed out of the error term for the equations in levels. Appendix B contains a detailed discussion of the assumptions required for the estimates to be consistent and discusses how these assumptions are tested in the data. The conclusion is that the empirical model appears to be well-specified in responding to concerns about serial correlation and potential endogeneity of the variables of interest.

2.6. Results: Firm-level determinants of the labour share of income

This section presents the findings on the relationship between market power, capital intensity, and the labour share. I begin by showing the results of estimating static models of the labour share to show that the results suffer from a pattern of serial correlation in the residuals that suggests a need for a dynamic model. I then present the results of estimating a dynamic model and explore the robustness of the results to alternative sample and variable definitions, and empirical specifications.

Table 2.3 presents the results of estimating versions of the static model given by equation (2.4). The dependent variable in all columns is the log labour share, defined as total compensation divided by value added. Column 1 regresses the log labour share on capital intensity, measured as the log of net tangible assets per employee, and market share, measured as the log of the firm's share of 4-digit sector sales. The elasticity of the labour share with respect to capital intensity is estimated to be -0.079. In other words, a doubling

¹¹ GMM estimation is carried out using the *xtabond2* command in Stata, implemented by Roodman (2009), with Windmeijer's (2005) finite-sample correction to the standard errors. I also report OLS results with and without firm fixed effects, to provide context for the GMM results.

of capital intensity is associated with a 7.9% reduction in the labour share.¹² The estimated elasticity with respect to market share is much smaller in magnitude, at -0.009. A doubling of the market share is thus associated with a 0.9% fall in the labour share. However, both effects are highly statistically significant, with a p-value below 1%. Column 2 adds sector fixed effects to control for sector-level differences in the labour share. The capital intensity estimate falls slightly to -0.067, while the market share estimate increases substantially in magnitude to -0.020. Both estimates remain highly statistically significant. Column 3 adds region fixed effects to control for any geographic differences in labour share setting that could be driven, for instance, by labour market conditions, but the impact on the results is negligible.

OLS estimates thus highlight strong and significant correlations between firms' labour shares, market shares, and the capital intensity of production. However, as discussed earlier, the coefficients will be biased in the presence of unobserved heterogeneity in how the labour share is determined at the firm level. Column 4 includes firm fixed effects as one way to account for such unobserved heterogeneity. Firm-specific drivers of the labour share may include things like within-sector product differentiation, differences in production costs relative to competitors, or other aspects of the firm's management and organizational structure. The results with fixed effects represent within-firm estimates of the effects of an increase in capital intensity or market share on the labour share. They continue to show a strong negative relationship between the labour share and both market power and capital intensity.

As discussed in Section 2.5, fixed-effects estimates are biased if productivity shocks are serially correlated. The Arellano and Bond (1991) tests for serial correlation shown in

¹² Note that this is a percentage impact, not a percentage-point impact. The estimated elasticities are interpreted in a more intuitive way further below, after the presentation of the results.

Table 2.3 cast doubt on the validity of the static model. Table 2.4, therefore, presents the results of estimating the dynamic model of the labour share given by equation (2.5). Columns 1 and 2 show the estimates using OLS and fixed-effects (within-firm) estimators. The long-run elasticities are relatively similar to the estimates in the static model. However, the residuals continue to suffer from serial correlation.

Column 3 presents the results of GMM estimation, which uses instrumental variable methods to address potential sources of correlation between the error term and the regressors. The elasticity of the labour share with respect to capital intensity is -0.077, similar to the OLS results, while the elasticity with respect to market share is -0.184, relatively similar to the results with firm fixed effects. Both coefficients are highly statistically significant, suggesting strong links between firms' labour shares, market shares, and capital intensity. The test for serial correlation does not highlight any second-order serial correlation in the residuals (p-value 0.903), implying that there is no need to add more lags of the model variables to ensure consistent estimation in the GMM framework. Moreover, the Hansen test of overidentifying restrictions does not reject the null hypothesis of instrument exogeneity (p-value 0.459). Finally, the validity of the GMM approach is supported by the observation that the coefficient on the lagged dependent variable is between the OLS and fixed-effects estimates (Blundell and Bond, 2000). Together with the results of further tests in Appendix B, these diagnostics suggest that the model is appropriate for the data and is estimated consistently.

How robust are the results to the way the sample has been defined and to the details of the empirical specification? This is tested in Table 2.5 by modifying the baseline GMM specification described above in several ways. Columns 1 to 3 address potential concerns with sample selection. One such concern is that the universe of incorporated firms that submit financial reports to Companies House may include self-employed individuals. I

therefore check whether the estimated coefficients change when firms that are never observed with more than one employees are excluded from the sample. Column 1 reports these results. The estimated elasticities are now -0.070 in the case of capital intensity and -0.192 in the case of the market share, both significant at a 1% level and similar in magnitude to those obtained from the full sample. The main reason is that the smallest firms are typically not required to supply fully detailed financial statements, and therefore few of them make it into the main sample. Next, column 2 focuses on firms observed in each year of the sample period. These tend to be larger and more established firms. This sub-sample is therefore more likely to satisfy the assumptions of System GMM (see Appendix B for a discussion). Again, the results are similar to the baseline model in column 5 of Table 2.3. The estimated effect of capital intensity on the labour share is larger in magnitude at -0.095, but the difference is not significant, while the estimated effect of the market share is almost identical at -0.182. Finally, column 3 tests the impact of outliers by estimating the model using unwinsorised variables. While this adds some noise to the data and the coefficient estimates shift slightly, the results remain very similar to those in the baseline GMM model.

Columns (4) to (7) explore different functional forms. Column 4 presents estimates in levels, rather than logs. In column 5 only capital intensity is in logs, while the labour share and market share are in levels. In column 6, only the labour share is in levels. Column 7 adds the square of capital intensity to the log-log model, motivated by the non-linear relationship found in plots of the raw data. The market share remains negative and statistically significant across these specifications, with differing magnitudes in columns 4 to 6 reflecting the fact that the interpretation of the coefficient changes depending on whether the variables are expressed in logs or in levels. Similarly, market share enters negatively and significantly when its square is included in the log-log model, with an

elasticity at the mean similar to that obtained in the linear model. Capital intensity is somewhat less robust to changes in the model specification.

Table 2.6 shows that the overall conclusions are also robust to several alternative definitions of the firm-level labour share, capital intensity, and market power. Column 1 replaces the dependent variable with the wage share of value added, excluding non-wage benefits such as pensions and social security taxes from the numerator. The estimated elasticities with respect to capital intensity and market share of -0.070 and -0.201, respectively, differ little from those in the baseline specification. This suggests that the effects on the labour share operate mainly through wages, and not through other forms of compensation.

Next, columns 2 to 5 test alternative measures of market power. A firm's share of 4-digit sector value added, rather than sales, is used in column 2, and the share of sales within the sector and NUTS3 region (approximately equivalent to a county) is used in column 3. Column 4 then repeats the specification from column 3 while restricting the sample to non-tradable sectors, as defined by Mano and Castillo (2015). Finally, column 5 estimates market power by the log of the ratio of value added to sales, as a proxy for the mark-up. Since this variable is highly correlated with the denominator of the labour share, it is instrumented with lags dated $t - 3$ and $t - 4$ for the equation in levels and $t - 4$ and $t - 5$ for the equation in first differences, to avoid an overlap with contemporaneous instruments for the lagged labour share. The results of all the analyses of alternative measures of market power are directionally the same as in the main specification in Table 2.4. In columns 2 to 4, the coefficients on market share range from -0.108 to -0.158 and the coefficients on capital intensity range from -0.068 to -0.079. In column 5, the estimated coefficient on the mark-up is -0.660, which suggests that for an average firm, a 1 percentage-point increase in the mark-up leads to a 1.7 percentage-point decrease in the

labour share. All results remain statistically significant, with the exception of the capital intensity variable in the last column. This confirms that the conclusions of the paper are robust to alternative measures of key variables.

Overall, a firm's market power and capital intensity has significant effects on the labour share. How do these results compare to those found in the macroeconomic literature? In terms of the market share, the work of Autor et al. (2017) and De Loecker and Eeckhout (2017), among others, finds a negative relationship between various measures of industry concentration and the industry-level labour share. My firm-level conclusions are consistent with aggregate findings, provided that firms' shares of sales and of value added do not move strongly in opposite directions.

The second finding that higher capital intensity lowers the labour share is consistent with the claims of Bentolila and Saint-Paul (2003), Karabarbounis and Neiman (2014) and Piketty (2014), who explain the trends in aggregate factor income shares through an elasticity of substitution between capital and labour greater than one. However, as mentioned earlier, the degree to which capital and labour are substitutable is a matter of ongoing debate in the literature.

One reason why the conclusions about the substitutability of capital and labour may vary is that different types of capital assets and workers may substitute or complement each other differently. I therefore probe the relationship between capital intensity and the labour share in three ways. First, to determine whether the effect of capital intensity differs according to the skill level of a firm's workforce, I split the sample into low-wage and high-wage firms. I define high-wage firms as those that, on average, are in the top half of the annual distribution of total compensation per employee during their presence in the sample. Given that firms' financial reports contain no information about the characteristics of

individual workers, I take the average wage level to be a proxy for the skill level of the employees.¹³ This helps judge whether the conclusions about the degree of substitutability between capital and labour vary depending on the type of workforce that firms employ. Second, to determine whether labour is equally substitutable with different classes of capital, I add intangible capital intensity, defined as the log of the book value of intangible assets per employee, as a regressor. Intangible assets reflect items such as patents, copyrights, and software, although firms are generally only able to recognise the fair value of acquired, rather than internally-generated, intangibles. Third, I combine the two approaches above to determine how the two types of workforces and two types of capital interact.

It turns out that the aggregate results on capital intensity are driven mainly by low-wage firms, suggesting that a low-skilled workforce is more substitutable with capital than a high-skilled workforce. Moreover, higher intangible capital intensity leads to a higher labour share, but only in the case of firms with a high-skilled workforce. These results are presented in Table 2.7. In column 1, tangible capital intensity is interacted with an indicator for whether the firm is a high-wage firm. The estimated elasticity for low-wage firms is -0.126, larger in magnitude than the aggregate estimate and highly significant. While the estimate for high-wage firms is also negative at -0.034, it is smaller and not statistically significant. Therefore, to the extent that a firm's average wage level can proxy for the average skills of its workforce, these results suggest that high-skilled workers are less substitutable with (tangible) capital than low-skilled workers, although the wide GMM standard errors mean that the difference between the two groups is not statistically

¹³ Table 2.C2 in the Appendix shows that the sample firms classified as high-wage are more concentrated in sectors such as manufacturing, construction, information and communication, and professional, scientific, and technical activities, where data from the UK Labour Force Survey shows that both a) average wages and b) the average skill level in the employee population (measured through years of education and the proportion of employees with a college degree) are indeed higher.

significant. Next, column 2 shows that the estimated elasticity of the labour share with respect to intangible capital intensity is 0.023 and statistically significantly different from the elasticity of the labour share with respect to tangible capital intensity. Column 3 shows that the estimate on intangible capital intensity is driven entirely by high-wage firms, although again, the size of the standard errors makes the differences between some groups insignificant. On the assumption that firms with high average wages have a workforce with high average skills, this suggests that high-skilled workers may be less substitutable with intangible capital than low-skilled workers, while low-skilled workers may be more substitutable with tangible capital than high-skilled workers. These conclusions would be consistent with the hypothesis of capital-skill complementarity (Griliches, 1969) and the findings of the literature on skill-biased technical change (see Hornstein, Krusell, and Violante, 2005, for a review).

How can we interpret the magnitudes of the estimated elasticities? One way to answer this question is to note that the aggregate labour share in the G7 countries has declined by approximately 10 percentage points since 1970, and to ask what size increase in the market share or capital intensity of the average firm would be required to obtain a reduction in the labour share of this magnitude. Using the estimates from column 3 of Table 2.7, the labour share of the average firm will decline by 10 percentage points if market share rises by 1.2 percentage point, holding capital intensity constant. On the other hand, if market share is held constant, then the same 10 percentage-point decline in the labour share will arise if capital intensity rises by approximately £100,000 for high-wage firms and £42,000 for low-wage firms. Another way to think about these amounts is to compare them with the median real employee compensation cost of £29,500. Thus, net investment equal to 1.4 years of labour costs at low-wage firms and 3.4 years of labour costs at high-wage

firms would be needed to reduce the labour share of income by 10 points, all else remaining constant.

An alternative way to assess the magnitude of the estimates in the context of the sample is to consider the actual dispersion of market shares and capital intensities across firms. Most firms observed in the data have low market shares and low capital intensity, but the distributions of both variables are highly skewed to the right, with a large standard deviation relative to the mean. A one standard deviation rise in market share (4.4 percentage points) is estimated to lower the labour share of the average firm by 36.8 percentage points. A move of this size implies almost a tripling of the market share by the average firm. On the other hand, a one standard deviation rise in net tangible assets per employee reduces the labour share of the average firm by 11.7 percentage points in the case of high-wage firms and 19.3 percentage points in the case of low-wage firms. While large, the effect is relatively smaller than that for market share, given the smaller absolute size of the estimated elasticity of the labour share with respect to capital intensity.

2.7. Human capital, labour market power, and the labour share

While the results for high- and low-wage firms shed light on the possible mechanisms through which capital intensity impacts the labour share, how a firm's income is shared may depend also on the relative bargaining power of the firm and its workers. In this section, I examine the impact of workforce characteristics and the firm's labour market power on the labour share in two ways. First, I explore whether enhancing the empirical model to include human capital variables changes the main findings. Second, I test two ways to account explicitly for the effect of the firm's labour market power (as distinct from product market power) on the labour share. Additionally, I use the information on average

workforce characteristics to address potential measurement problems with the capital intensity variable.

One aspect of labour share determination that the results in the previous section do not take into account is the firm's position in the labour market. So far, I have assumed that labour is supplied elastically to the firm at a constant wage w . To the extent that labour markets are not perfectly competitive, variations in the labour share—both within and across firms—may be driven by differences in labour market power. Such differences may arise, for instance, if workers are not perfectly mobile across geographies. Moreover, even firms in the same sector and location may demand workers with different skills and thus participate in separate labour markets (at least in the short term). This can happen, for example, if unobserved product differentiation leads to heterogeneous complementarities between firms' capital assets and different classes of labour. In contrast, unionization, or other forms of collective bargaining, may enhance the bargaining power of workers.

The firm's position in the labour market may be an omitted variable affecting the results for product market power and capital intensity. This is because labour market characteristics might be correlated with firms' current or historical capital investments and market position. For instance, the tightness of the labour market in a specific skill category might affect the type of capital that the firm needs to employ in production and the national market share that it can achieve, or vice versa.

Modifying the simple model outlined in Section 2.3 to allow for imperfect competition in the labour market predicts that the labour share will be inversely related to the firm's labour market power. The first-order condition for labour becomes:

$$s_i^L = \frac{\varepsilon_i^{YL} \left(1 + \frac{1}{\eta}\right)}{1 + \frac{1}{\lambda}} \quad (2.6)$$

where $\lambda \geq 0$ is the labour supply elasticity. Inelastic supply of labour to the firm (low λ) reduces the labour share compared to the case of perfectly competitive labour markets (infinitely high λ). This is because highly inelastic labour supply gives the firm greater ability to reduce wages without materially restricting the supply of workers willing to work there. This situation may arise if there are frictions in the labour market, such as mobility costs that limit workers' ability to seek alternative employment.¹⁴ A firm's labour market power may thus be a relevant variable to take into account when estimating the determinants of the labour share.

Unfortunately, FAME does not contain any information about firms' workforce other than total employment. My first step, therefore, is to use the UK Labour Force Survey to construct a range of human capital variables that describe the average workforce characteristics in each 1-digit sector-region-year cell. These human capital variables mirror standard variables found in wage regressions: education, experience, and the percentage of employees that are female, married, non-white, UK-born, and that work part-time. I add them to the model and check whether a) cells with a greater proportion of groups that might be expected to have lower bargaining power in the labour market are associated with a lower labour share, and b) whether controlling for these workforce characteristics changes the earlier results on the role of capital intensity and market share. For instance, immigrants, part-time workers, or other groups with weaker attachment to the labour force or the firm may have little influence on rent-sharing policies.

¹⁴ Manning (2003) sets out a comprehensive case for the empirical relevance of imperfect competition in the labour market, i.e., the supply of labour to an individual firm not being infinitely elastic.

Column 1 of Table 2.8 shows that the estimates for the market share and capital intensity are unaffected when the human capital variables are simply added as controls to the main specification. The labour share appears to be, on average, lower in sector-region cells where a higher proportion of employees work part time or hold “other qualifications,” which corresponds to low-level vocational and non-standard qualifications. A higher labour market participation rate is also associated with a lower average labour share. These results give some—albeit imperfect—indication that the labour share is lower at firms operating in sectors and regions where the average worker is more weakly positioned in the labour market. However, these coefficients have many possible interpretations. For instance, the coefficient on the percentage of college-educated employees is negative and significant. Perhaps some of the occupations in this category are susceptible to outsourcing or replacement by machines, giving firms greater power over wage setting, but the mechanism is unclear. The data cannot distinguish to what extent the estimated relationship between workforce characteristics is driven by the tension between the strength of a given group’s attachment to the labour market, with the consequential impact on bargaining power, social norms, and various forms of selection.

I therefore attempt to measure firms’ labour market power, over and above these average workforce characteristics, in two alternative ways. The first is the proportion of new recruits from non-employment in each sector-region cell. Manning (2003) suggests this as a measure of firms’ wage-setting power, on the basis that their ability to reduce wages is kept in check by workers’ ability to leave for other employers. The higher the proportion of recruits from non-employment, the lower the competition for workers among firms and more monopsonistic the labour market. Manning (2003) also shows that in a general equilibrium model of the labour market with search frictions, this ratio is monotonically inversely related to the ratio of the job arrival rate to the job destruction rate.

As it increases, each worker's wage approaches the marginal product. Thus, a high proportion of recruits from non-employment implies that firms have the power to pay workers significantly less than the marginal product and the labour market is farther away from perfect competition. Column 2 of Table 2.8 shows that labour market power measured this way is highly significant, with an estimated elasticity of -0.326. This implies that, at the mean, a one standard deviation increase in the proportion of recruits from non-employment (4 percentage points) is associated with a 1.6 percentage-point reduction in the labour share.¹⁵

Ideally, the proportion of recruits from non-employment would be calculated at the level of individual firms. I therefore use a second, alternative measure of a firm's labour market position: the log of the number of firms present in the same NUTS3 region (roughly a county) or in the same region-industry cell, as proxies for competition for labour in the local labour market. The results are shown in columns 3 and 4 of Table 2.8, respectively. To make the sign of the estimate comparable to the other columns, the number of competing firms is multiplied by -1. The estimated elasticity is -0.023. This suggests that a higher number of firms raises the labour share. Perhaps it is easier for workers in such markets to switch jobs without incurring substantial mobility costs, which improves their bargaining position.

There is one other way in which these human variables are helpful, and that is to check whether a potential measurement issue with the capital-labour ratio affects the conclusions of this paper. In the theoretical model outlined in Section 2.3, capital and labour are homogeneous, while in practice, there is likely to be considerable heterogeneity across

¹⁵ The mean proportion of new recruits from non-employment in the sample is 0.61. While this may seem high, it is in the range reported by Manning (2003) for the UK and the US using different periods and data sources.

firms in the quality of their capital and labour inputs. To the extent that differences in the quality of capital equipment and structures are reflected in the purchase price, they will be reflected in the book value of tangible assets—the numerator of the capital intensity measure. Heterogeneity in worker quality, however, will not be reflected in the number of employees, which forms the denominator of the ratio. This suggests a potential measurement problem related to the differences in the average wage across firms that may also be related to the labour share. One solution would be to measure the labour input using compensation cost, on the assumption that wage differences fully reflect the differences in worker quality. However, the denominator of the capital intensity would then be the same as the numerator of the dependent variable.

My approach is therefore to use the human capital variables to derive a quality adjustment for each firm's labour input and use this estimate of "effective labour" to calculate an alternative capital intensity measure. I begin by regressing the average compensation per employee at each firm on the vector of human capital variables. I use the estimated coefficients from this regression to predict an average wage for each firm. I then adjust each firm's total number of employees by the percentage deviation of its predicted wage from the mean of all of that year's predictions. A firm observed in a sector-region cell where the characteristics of the workforce predict relatively low wages will thus be treated as having a lower effective labour input, and thus higher quality-adjusted capital intensity. The resulting headcount adjustment varies from -43% to +47% and accounting for this increases the standard deviation of measured labour input by 10.2%, compared to using employment as the measure of labour input.

Column 5 of Table 2.8 shows the results. Since the human capital variables are used as instruments to derive a predicted labour input, they are no longer used as exogenous control variables. Replacing the standard capital intensity variable with the adjusted

measure gives similar estimates as the baseline specification in Table 2.3. If anything, the estimated elasticity with respect to the capital-labour ratio is slightly larger in magnitude (-0.084 vs. -0.077), while the estimate on market share is slightly smaller (-0.179 vs. -0.184), but the differences are not statistically significant. While the employment adjustment is highly imperfect due to the lack of firm-specific labour quality measures in the dataset, it does not highlight any obvious impact of systematic variation in labour quality on the conclusions.

Together, these results suggest that a firm's labour market position may be relevant for labour share determination. However, the dataset does not lend itself particularly well to deriving firm-specific variables that isolate this aspect of labour share setting. Therefore, further exploration of the impact of imperfect competition in the labour market on the labour share, using firm-specific data sources or matched employer-employee datasets, would be a valuable direction for future work—especially since a firm may wield different degrees of power in different labour markets where it recruits employees of different skill levels and occupations.¹⁶ Notably, however, the estimated effects of capital intensity and market share on the labour share do not change with the addition of any of the proposed human capital and labour market variables. They are also unaffected by the attempt to adjust firms' labour input for the average characteristics of the workforce in their sector and region. The main conclusions of this paper are still that firms with greater market power and higher capital intensity share a smaller proportion of their income with workers.

¹⁶ In the UK, available matched employer-employee data suffers from two shortcomings. First, firm information does not include a balance sheet, making it difficult to test the role of capital intensity in labour share determination. Second, employee data comes from a 1% nationwide sample, so it is not possible to characterize the within-firm distribution of employees or new hires. However, data from other countries may be helpful for shedding further light on these questions.

2.8. Conclusions

This paper considers the relationship between labour share of income at the level of an individual firm and two variables suggested by economic theory: market power and capital intensity of production. It contributes to the literature on factor income shares by documenting the wide dispersion of firm-level labour shares and by estimating a model of labour share determination using longitudinal firm data, with a focus on addressing endogeneity concerns through instrumental variable methods. Its main finding is that the capital intensity of firms' production and their power in the product market are significant determinants of the labour share, consistent with economic intuition and a simple theoretical model. The labour share is lower when firms employ more capital relative to labour. This result is driven by low-wage firms, suggesting that the elasticity of substitution between capital and labour may be greater than one at firms that employ a low-skilled workforce. In addition, the labour share is lower when firms have greater power in the product market, measured by their share of sales or value added. This is also consistent with intuition: firms that face less elastic demand raise their profit share of output to the detriment of the labour share. GMM estimation suggests further that the effect of the market share is quantitatively larger than that of capital intensity.

The main implication of this paper is that the aggregate labour share of income is impacted to a large extent by the decisions made at the level of individual firms, as well as by the environment that those firms face. To the extent that the relationships estimated here hold in the future, increasing robotization may augur continued declines for the labour share, particularly in the segments of the economy characterized by low-wage employment. Similarly, the labour share may continue to be under pressure in markets where a few large firms command significant or rising pricing power.

Future research could focus on relating these results back to macroeconomic trends by decomposing the historical movements in the aggregate labour share into changes at the firm level and changes in the distribution of firms across capital intensity and market share categories within sectors. The analysis could also be replicated using firm-level data from other countries and expanded to longer time periods that coincide with more dramatic labour share declines. Finally, causal links between capital intensity, market power, and the labour share can also be explored further, perhaps by identifying natural experiments that provide quasi-exogenous variation in access to capital across firms or in market structure across sectors.

2.9. Bibliography


- Acemoglu, D. (2003). 'Labor- and capital-augmenting technical change', *Journal of the European Economic Association*, vol. 1(1), pp. 1-37.
- Arellano, M. and Bond, S. (1991). 'Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations', *Review of Economic Studies*, vol. 58(2), pp. 277-97.
- Arellano, M. and Bover, O. (1995). 'Another look at the instrumental variable estimation of error-components models', *Journal of Econometrics*, vol. 68(1), pp. 29-51.
- Arpaia, A., Pérez, E. and Pichelmann, K. (2009). 'Understanding labour income share dynamics in Europe', MPRA Paper 15649.
- Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2017) 'The fall of the labor share and the rise of superstar firms', NBER Working Paper No. 23396.
- Azmat, G., Manning, A. and Van Reenen, J. (2012). 'Privatization and the decline of labour's share: international evidence from network industries', *Economica*, vol. 79(315), pp. 470-92.
- Barkai, S. (2016). 'Declining labor and capital shares', Working Paper, University of Chicago.
- Bentolila, S. and Saint-Paul, G. (2003). 'Explaining movements in the labor share', *Contributions to Macroeconomics*, vol. 3(1), pp. 1-33.
- Blanchard, O. and Giavazzi, F. (2003). 'Macroeconomic effects of regulation and deregulation in goods and labor markets', *Quarterly Journal of Economics*, vol. 118(3), pp. 879-907.
- Blundell, R. and Bond, S. (1998). 'Initial conditions and moment restrictions in dynamic panel data models', *Journal of Econometrics*, vol. 87(1), pp. 115-43.
- Blundell, R. and Bond, S. (2000). 'GMM estimation with persistent panel data: an application to production functions', *Econometric Reviews*, vol. 19(3), pp. 321-40.
- Bond, S. and Guceri, I. (2017). 'R&D and productivity: evidence from large UK establishments with substantial R&D activities', *Economics of Innovation and New Technology*, vol. 26(1-2), pp. 108-120.
- Bond, S. and Söderbom, M. (2005). 'Adjustment costs and the identification of Cobb-Douglas production functions', IFS Working Paper 05/04.
- Böckerman, P. and Maliranta, M. (2012). 'Globalization, creative destruction, and the labour share change: evidence on the determinants and mechanisms from longitudinal plant-level data', *Oxford Economic Papers*, vol. 64(2), pp. 259-280.

- Chirinko, R. (2008). 'σ: The long and short of it', CESifo Working Paper 2234.
- Companies House. (2015). *Annual Report and Accounts 2014/15*. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/446159/Annual_Report_201415.pdf (last accessed 30 December 2017).
- De Loecker, J. and Eeckhout, J. (2017). 'The rise of market power and the macroeconomic implications', NBER Working Paper No. 23687.
- Draca, M., Machin, S., and Van Reenen, J. (2011). 'Minimum wages and firm profitability', *American Economic Journal: Applied Economics*, vol.3 (1), pp.129-151.
- Elsby, M., Hobijn, B. and Şahin, A. (2013). 'The decline of the U.S. labor share', *Brookings Papers on Economic Activity*, vol. 2013(2), pp. 1-63.
- Gollin, D. (2002). 'Getting income shares right', *Journal of Political Economy*, vol. 110(2), pp. 458-74.
- Griliches, Z. (1969). 'Capital-skill complementarity', *Review of Economics and Statistics*, vol. 51(4), pp. 465-68.
- Growiec, J. (2012). 'Determinants of the labor share: evidence from a panel of firms', *Eastern European Economics*, vol. 50(5), pp. 23-65.
- Guerriero, M. (2012). 'The labour share of income around the world. Evidence from a panel dataset', University of Manchester Working Paper No. 32/2012.
- Haskel, J. and Westlake, S. (2018). *Capitalism without capital: The rise of the intangible economy*, Princeton, NJ: Princeton University Press.
- Hicks, J. (1932). *The theory of wages*, London: Macmillan.
- Hornstein, A., Krusell, P., and Violante, G. (2005). 'The effects of technical change on labor market inequalities' in: *Handbook of Economic Growth*, Vol.1, pp.1275-1370.
- IMF. (2017). 'Understanding the downward trend in labor income shares', in *World Economic Outlook, April 2017*.
- Kaldor, N. (1957). 'A model of economic growth,' *Economic Journal*, vol. 67, pp. 591-624.
- Karabarbounis, L. and Neiman, B. (2014). 'The global decline of the labor share', *Quarterly Journal of Economics*, vol. 129(1), pp. 61-103.
- Keynes, J. M. (1939). 'Relative movements in real wages and output,' *Economic Journal*, vol. 49, pp. 34-51.
- Knöblach, M., Rößler, M., and Zwerschke, P. (2016). 'The elasticity of factor substitution between capital and labour in the U.S. economy: a meta-regression analysis', CEPIE Working Paper 03/16.

- Manning, A. (2003). *Monopsony in motion*, Princeton, NJ: Princeton University Press.
- Mano, R. and Castillo, M. (2015). ‘The level of productivity in traded and non-traded sectors for a large panel of countries’, IMF Working Paper 15/48.
- ONS. (2016). *UK National Accounts, The Blue Book: 2016*. London: Office for National Statistics.
- ONS. (2018). *UK Non-Financial Business Economy (Annual Business Survey)*. <https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/datasets/uknonfinancialbusinesseconomyannualbusinesssurveysectionsas> (last accessed: 11 April 2018).
- Pessoa, J. P. and Van Reenen, J. (2013). ‘Decoupling of wage growth and productivity growth: myth and reality’, CEP Discussion Paper 1246.
- Piketty, T. (2014). *Capital in the twenty-first century*, Cambridge, MA: The Belknap Press of Harvard University Press.
- Robinson, J. (1933). *The economics of perfect competition*, London: Macmillan.
- Rodríguez, F. and Jadayev, A. (2010). ‘The declining labor share of income’, UNDP Human Development Research Paper 2010/36.
- Roodman, D. (2009). ‘How to do xtabond2: an introduction to difference and system GMM in Stata’, *Stata Journal*, vol. 9(1), pp. 86-136.
- Siegenthaler, M. and Stucki, T. (2015). ‘Dividing the pie: the determinants of labor's share of income on the firm level’, *ILR Review*, vol. 68(5), pp. 1157-1194.
- Windmeijer, F. (2005). ‘A finite sample correction for the variance of linear efficient two-step GMM estimators’, *Journal of Econometrics*, vol. 126(1), pp. 25-51.


2.10. Figures and tables

FIGURE 2.1: LABOUR SHARE TRENDS SINCE 1970




Aggregate labour share is defined as compensation of employees divided by GVA minus mixed income to exclude the output of the self-employed from the denominator (see Gollin, 2002). Data for G7 countries excluding the UK (Canada, France, Germany, Italy, Japan, US) is weighted by GVA. Germany prior to 1990 reflects West Germany only. Source: OECD, own calculations.

FIGURE 2.2: LABOUR SHARE TRENDS OVER THE SAMPLE PERIOD, 2005-2012



The labour share for the whole UK economy and for G7 countries excluding the UK is calculated as aggregate compensation of employees divided by GVA minus mixed income. Data for the UK non-financial business sector is based on published ONS Annual Business Survey results. Dashed lines represent the lines of best fit. The individual plots are on different scales to enable a side-by-side comparison of cyclical fluctuations.

FIGURE 2.3: LABOUR SHARE VS. CAPITAL INTENSITY AND MARKET SHARE



Local polynomial plots of the log labour share against the capital intensity and market share variables for the GMM estimation sample from column 5 of Table 3. The shaded areas correspond to 95% confidence intervals.

TABLE 2.1: LABOUR SHARE DISPERSION

Variable, x	Between firms $\sigma_u/ \bar{x} $	Within firms $\sigma_\varepsilon/ \bar{x} $	Between % ρ
Labour Share	0.40	0.60	0.31
Log(Labour Share)	1.14	1.08	0.53
Average Compensation	0.41	0.18	0.84
Log(Average Compensation)	0.13	0.05	0.88
Average Product of Labour	0.88	0.54	0.73
Log(Average Product of Labour)	0.15	0.11	0.67

This table presents dispersion metrics for the labour share and its components, both in levels and in logs. The values σ_u and σ_ε are estimated standard deviations of the error terms from the random-effects model, $x_{it} = \alpha + u_i + \varepsilon_{it}$, normalized by the means to facilitate comparisons across variables. The last column shows the proportion of variance due to differences across firms, $\rho = \sigma_u^2/(\sigma_u^2 + \sigma_\varepsilon^2)$. The model includes year and 5-digit sector fixed effects to account for aggregate trends and cross-sector differences in the labour share. Sample size is 119,764 observations.

TABLE 2.2: SUMMARY STATISTICS

	Observations	Mean	Median	Std. Dev.	Min	Max
Sales (£ 000)	119,764	67,006	12,105	633,838	2	101,199,900
Employees	119,764	317	73	2,926	1	277,684
Net Tangible Assets (£ 000)	119,764	17,123	974	228,054	1	19,407,766
Labour share (Compensation / Value Added)	119,764	0.83	0.75	0.61	0.00	5.01
Net Tangible Assets / Employee (£ 000)	119,764	40	11	102	0	796
Market Share in 4-Digit Sector	119,764	0.02	0.00	0.04	0.00	0.30
<i>Labour Share by Capital Intensity (Net Tangible Assets/Employee):</i>	Observations	Mean	Median	Median (Manufacturing)	Median (Services)	
<£10,000	56,415	0.90	0.81	0.80	0.80	
£10,000-50,000	44,186	0.82	0.74	0.74	0.74	
£50,000-100,000	10,204	0.77	0.66	0.64	0.68	
£100,000-500,000	7,144	0.63	0.52	0.54	0.52	
≥£500,000	1,815	0.35	0.20	0.27	0.24	
<i>Labour Share by Market Share in 4-Digit Sector:</i>	Observations	Mean	Median	Median (Manufacturing)	Median (Services)	
<1%	91,143	0.85	0.76	0.76	0.76	
1-5%	19,379	0.78	0.72	0.72	0.71	
5-10%	4,397	0.77	0.70	0.71	0.69	
10-20%	2,507	0.75	0.69	0.69	0.69	
≥20%	2,338	0.72	0.65	0.65	0.67	

Data for firm-level observations pooled across all years in the dataset. Sales and net tangible assets are measured in thousands of 2011 pounds. All ratios (compensation cost/employee, labour share, net tangible assets/employee, and market share) are winsorised at the 1st and 99th percentile.

TABLE 2.3: STATIC MODELS OF THE LABOUR SHARE

	OLS			Within-Firm
	(1)	(2)	(3)	(4)
Capital Intensity	-0.079*** (0.002)	-0.067*** (0.002)	-0.068*** (0.002)	-0.030*** (0.004)
Market Share	-0.009*** (0.001)	-0.020*** (0.002)	-0.019*** (0.002)	-0.219*** (0.008)
Serial Correlation Test				
1 st Order	0.000	0.000	0.000	0.000
2 nd Order	0.000	0.000	0.000	0.000
Year FE	x	x	x	x
Sector FE		x	x	
Region FE			x	
Firm FE				x
Sample Size	119,764	119,764	119,764	119,764

This table presents the results of estimating a static model of the log labour share. OLS results from estimating versions of equation (2.4) with varying fixed effects are presented in columns 1-3. Column 4 additionally includes firm fixed effects. Serial correlation test: p-value for the null hypothesis of no serial correlation from the Arellano and Bond (1991) test for first and second-order serial correlation. Robust standard errors clustered by firm are in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

TABLE 2.4: DYNAMIC MODELS OF THE LABOUR SHARE

	OLS	FE	GMM
	(1)	(2)	(3)
Labour Share (t-1)	0.549*** (0.005)	-0.029*** (0.006)	0.159*** (0.010)
Capital Intensity (t)	-0.050*** (0.004)	-0.035*** (0.004)	-0.476*** (0.142)
Capital Intensity (t-1)	0.016*** (0.004)	0.011*** (0.004)	0.411*** (0.127)
Market Share (t)	-0.278*** (0.009)	-0.304*** (0.009)	-0.057 (0.104)
Market Share (t-1)	0.274*** (0.009)	0.159*** (0.008)	-0.097 (0.091)
Long-Run Elasticities:			
Capital Intensity	-0.075*** (0.003)	0.024*** (0.004)	-0.077*** (0.022)
Market Share	-0.008*** (0.002)	-0.141*** (0.009)	-0.184*** (0.024)
Serial Correlation Test			
1 st Order	0.000	0.000	0.000
2 nd Order	0.000	0.000	0.903
Overidentification Test			0.459
Year FE	x	x	x
Sector FE	x		x
Region FE	x		x
Firm FE		x	
Sample Size	119,764	119,764	119,764

This table presents the main results of dynamic regressions of the log labour share on capital intensity, market share, and controls. Column 1 presents the OLS results from estimating the model

$$s_{it}^L = \rho s_{i,t-1}^L + \pi_1 \left(\frac{K_{it}}{L_{it}} \right) + \pi_2 \left(\frac{K_{i,t-1}}{L_{i,t-1}} \right) + \pi_3 s_{it}^M + \pi_4 s_{i,t-1}^M + (a_i + v_{it}).$$

Column 2 additionally includes firm fixed effects. Column 3 presents GMM results described in the main text, with the lagged levels of the labour share, capital intensity, and market share dated t-2 and t-3 used as instruments for the difference equation, and lagged differences dated t-1 and t-2 used as instruments for the levels equation. Serial correlation test: p-value for the null hypothesis of no serial correlation from the Arellano and Bond (1991) test for first and second-order serial correlation. Overidentification test: p-value for the null hypothesis of instrument exogeneity from the Hansen test of overidentifying restrictions. Robust standard errors clustered by firm are in parentheses, with Windmeijer's (2005) finite-sample correction applied to GMM. *** p<0.01, ** p<0.05, * p<0.10.

TABLE 2.5: ROBUSTNESS TO ALTERNATIVE EMPIRICAL SPECIFICATIONS

	Larger Firms	Balanced Sample	Including Outliers	Levels vs. Logs			Squared K/L
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Capital Intensity	-0.070*** (0.022)	-0.095** (0.040)	-0.056*** (0.020)	-0.000 (0.000)	-0.019 (0.029)	-0.005 (0.038)	-0.042* (0.022)
Market Share	-0.192*** (0.025)	-0.182*** (0.050)	-0.204*** (0.025)	-3.292** (1.370)	-3.602** (1.544)	-0.094*** (0.035)	-0.185*** (0.024)
Serial Correlation Test							
1 st Order	0.000	0.001	0.001	0.000	0.000	0.000	0.000
2 nd Order	0.971	0.792	0.919	0.930	0.949	0.950	0.658
Overidentification Test	0.416	0.981	0.619	0.276	0.232	0.115	0.588
Year FE	x	x	x	x	x	x	x
Sector FE	x	x	x	x	x	x	x
Region FE	x	x	x	x	x	x	x
Sample Size	119,121	26,878	119,764	129,878	129,065	129,052	119,764

This table presents alternative empirical specifications of the model relating the labour share to capital intensity and market share. The first three columns consider different definitions of the sample: (1) excluding firms that are only ever observed with one employee (potentially self-employed individuals operating as incorporated companies); (2) balanced sample of firms observed in all periods; (3) without winsorising any variables. Column 4 presents the results of estimating the model in levels instead of logs. In column 5, the labour share and market share are in levels while capital intensity is in logs, and in column 6, only the labour share is in levels. In column 7, the square of capital intensity is included and the reported coefficient is the elasticity calculated at the sample mean. Coefficients represent long-run elasticities. Serial correlation test: p-value for the null hypothesis of no serial correlation from the Arellano and Bond (1991) test for first and second-order serial correlation. Overidentification test: p-value for the null hypothesis of instrument exogeneity from the Hansen test of overidentifying restrictions. Robust standard errors clustered by firm are in parentheses, with Windmeijer's (2005) finite-sample correction applied to GMM. *** p<0.01, ** p<0.05, * p<0.10.

TABLE 2.6: ALTERNATIVE MEASURES OF THE LABOUR SHARE, CAPITAL INTENSITY, AND PRODUCT MARKET POWER

	Wage Share	Share of Sector Value Added	Market Share in Sector and Region	Non-Tradable Sectors	Mark-up
	(1)	(2)	(3)	(4)	(5)
Capital Intensity	-0.070*** (0.023)	-0.068*** (0.021)	-0.079*** (0.022)	-0.076** (0.037)	-0.045 (0.030)
Product Market Power	-0.201*** (0.028)	-0.154*** (0.030)	-0.158*** (0.020)	-0.108*** (0.027)	-0.660*** (0.091)
Serial Correlation Test					
1 st Order	0.000	0.000	0.000	0.000	0.000
2 nd Order	0.685	0.458	0.616	0.076	0.315
Overidentification Test	0.352	0.253	0.521	0.692	0.626
Year FE	x	x	x	x	x
Sector FE	x	x	x	x	x
Region FE	x	x	x	x	x
Sample Size	115,977	119,764	119,764	52,564	119,764

This table presents variations of the main GMM specification in column 3 of Table 2.4 using different definitions of the main variables of interest. In column 1, the dependent variable is the log of the wage share, defined as wages excluding non-wage compensation divided by value added. In column 2, market share is calculated as the share of total value added within a 4-digit sector. In columns 3 and 4, market share is calculated as sales share within a 4-digit sector and NUTS3 region (approximately a county), with the sample in column 4 limited to firms in non-tradable sectors. In column 5, the market share is replaced by the mark-up, calculated as the log of the ratio of value added before employment expense to revenues, as an alternative measure of market power. Coefficients represent long-run elasticities. Serial correlation test: p-value for the null hypothesis of no serial correlation from the Arellano and Bond (1991) test for first and second-order serial correlation. Overidentification test: p-value for the null hypothesis of instrument exogeneity from the Hansen test of overidentifying restrictions. Robust standard errors clustered by firm are in parentheses, with Windmeijer's (2005) finite-sample correction applied to GMM. *** p<0.01, ** p<0.05, * p<0.10.

TABLE 2.7: UNDERSTANDING HOW CAPITAL INTENSITY IMPACTS THE LABOUR SHARE

	(1)	(2)	(3)
Tangible Capital Intensity (Low-Wage Firms)	-0.126** (0.053)		-0.096** (0.043)
Tangible Capital Intensity (High-Wage Firms)	-0.034 (0.032)		-0.054* (0.029)
Tangible Capital Intensity		-0.078*** (0.019)	
Intangible Capital Intensity		0.023 (0.014)	
Intangible Capital Intensity (Low-Wage Firms)			-0.001 (0.016)
Intangible Capital Intensity (High-Wage Firms)			0.039* (0.023)
Market Share	-0.190*** (0.027)	-0.158*** (0.024)	-0.161*** (0.023)
Serial Correlation Test			
1 st Order	0.000	0.000	0.000
2 nd Order	0.576	0.706	0.833
Overidentification Test	0.444	0.400	0.433
Sample Size	119,764	118,500	118,500

This table extends the main GMM results from column 3 of Table 2.4 to further explore the relationship between capital intensity and the labour share. Tangible capital intensity is the log of the book value of tangible assets per employee (the same variable as capital intensity in Table 2.4). Intangible capital intensity is the log of the book value of intangible assets per employee. The high-wage firm indicator equals 1 if the firm is, on average, in the top half of the annual wage distribution during its presence in the sample, as a proxy for the average skill level of its workforce. Intangible capital All columns include year, sector, and region fixed effects. Coefficients represent long-run elasticities. Serial correlation test: p-value for the null hypothesis of no serial correlation from the Arellano and Bond (1991) test for first and second-order serial correlation. Overidentification test: p-value for the null hypothesis of instrument exogeneity from the Hansen test of overidentifying restrictions. Robust standard errors clustered by firm are in parentheses, with Windmeijer's (2005) finite-sample correction applied to GMM. *** p<0.01, ** p<0.05, * p<0.10.

TABLE 2.8: HUMAN CAPITAL, LABOUR MARKET POWER, AND THE LABOUR SHARE

	Human Capital	Recruits from Non-Employment	Employer Count		Effective Labour
	(1)	(2)	Region	Region-Industry	(5)
Capital Intensity	-0.078*** (0.022)	-0.078*** (0.022)	-0.080*** (0.022)	-0.079*** (0.022)	-0.084*** (0.023)
Market Share	-0.184*** (0.024)	-0.184*** (0.024)	-0.182*** (0.024)	-0.183*** (0.024)	-0.179*** (0.025)
Labour Market Power		-0.326*** (0.102)	-0.023*** (0.006)	-0.015*** (0.005)	
College %	-0.255*** (0.111)	-0.272*** (0.111)	-0.270** (0.112)	-0.266** (0.112)	
Other Higher Ed. %	-0.149 (0.151)	-0.134 (0.152)	-0.181 (0.151)	-0.183 (0.152)	
High School %	0.079 (0.124)	0.071 (0.125)	0.069 (0.125)	0.072 (0.125)	
Low Qualifications %	-0.102 (0.141)	-0.090 (0.142)	-0.117 (0.143)	-0.122 (0.143)	
Other Qualifications %	-0.328*** (0.135)	-0.347** (0.137)	-0.354** (0.137)	-0.343** (0.137)	
No Qualifications %	-0.126 (0.151)	-0.128 (0.151)	-0.130 (0.152)	-0.133 (0.153)	
Avg Experience	-0.022 (0.014)	-0.020 (0.014)	-0.021 (0.014)	-0.018 (0.014)	
Avg Experience ²	0.001 (0.000)	0.000 (0.000)	0.001 (0.000)	0.000 (0.000)	
Female %	0.135 (0.085)	0.130 (0.086)	0.144* (0.086)	0.158* (0.086)	
Non-White %	-0.092 (0.107)	-0.112 (0.108)	-0.073 (0.108)	-0.075 (0.108)	
UK-Born %	0.110 (0.103)	0.130 (0.103)	0.114 (0.103)	0.130 (0.104)	
Part-Time %	-0.600*** (0.107)	-0.612*** (0.108)	-0.604*** (0.107)	-0.590*** (0.107)	
Participation Rate	-0.455* (0.274)	-0.552** (0.279)	-0.433 (0.274)	-0.437 (0.274)	
Serial Correlation Test					
1 st Order	0.000	0.000	0.000	0.000	0.000
2 nd Order	0.912	0.957	0.910	0.907	0.825
Overidentification Test	0.500	0.488	0.506	0.511	0.493
Sample Size	119,764	119,764	119,362	119,362	119,764

This table adds human capital controls and measures of labour market power to the main specification from column 3 of Table 2.4. Coefficients represent long-run elasticities from GMM estimation. All columns include year, sector, and region fixed effects. Columns 1-4 also include human capital controls, defined as: the proportion of employees that are female, non-white, UK-born, married, part-time, and with different qualification levels; average years of education, experience, experience squared; and the participation rate – each calculated within a 1-digit sector, NUTS1 region, and year using the Labour Force Survey. Labour market power is defined as the proportion of recruits from non-employment in a sector-region cell in column 2; log of the total number of firms in the same NUTS3 region times -1 in column 3; and log of the total number of firms in the same NUTS3 region and 2-digit sector times -1 in column 4. Serial correlation test: p-value for the null hypothesis of no serial correlation from the Arellano and Bond (1991) test for 1st and 2nd-order serial correlation. Overidentification test: p-value for the null hypothesis of instrument exogeneity from the Hansen test of overidentifying restrictions. Robust standard errors are clustered by firm, with Windmeijer's (2005) finite-sample correction. *** p<0.01, ** p<0.05, * p<0.10.

2.11. Appendices

Appendix A: Time series properties of key variables

This Appendix provides evidence that the main variables of interest are persistent—a pattern that motivates the use of a dynamic model of the labour share in the empirical analysis. The analysis is similar to that carried out by Blundell and Bond (2000) in the context of estimating firm production functions.

Table 2.A1 reports simple AR(1) specifications of the labour share, capital intensity, and market share variables, estimated using the OLS, within-firm, difference GMM, and system GMM estimators. The reason for showing this range of estimators is that each of them may be biased under different conditions, but they can be used together to draw conclusions about the persistence of each series. For example, in the presence of firm-specific effects, OLS estimates of are expected to be biased upwards, while within-firm estimates are expected to be biased downwards. The OLS results in column 1 and the within-firm results in column 2 thus bracket the autoregressive coefficient. They suggest that capital intensity and market share are indeed highly persistent, while the labour share is somewhat less so.

GMM results help pin down the degree of persistence of each series. Column 3 shows difference GMM estimates, which model the process in differences and use lagged levels dated $t - 2$ or earlier as instruments for the lagged first difference of the dependent variable. These estimates suggest that while capital intensity is much more highly persistent than the labour share and the market share. However, difference GMM can be biased towards the within-firm estimator in finite samples when the instruments are weak.¹⁷

¹⁷ Table 2.A2 estimates reduced-form “first-stage” equations for the last year of the sample, when the greatest number of lagged values is available to serve as instruments. It shows that the instruments for the lagged difference of capital intensity and the market share in the first-differenced AR(1) model do indeed appear to be weaker than in the case of the labour share.

Column 4 therefore reports the results of the system GMM estimator, which can help address this bias under certain conditions (see Blundell and Bond, 2000). The system GMM specification reported here uses lagged levels dated $t - 2$ or earlier as instrumented in the first-differenced equation and the lagged difference dated $t - 1$ as an instrument in the levels equation. The results suggest that both capital intensity and the market share are highly persistent, with autoregressive coefficients greater than 0.9 and relatively similar to the OLS results in column 1, which served as an upper bound for the estimates. This persistence motivates the use of a dynamic specification in the main body of the paper.

TABLE 2.A1: AUTOREGRESSIVE MODELS OF THE MAIN VARIABLES OF INTEREST

	OLS (1)	Within Firms (2)	Difference GMM (3)	System GMM (4)
(a) Labour share				
Labour share (t-1)	0.598*** (0.005)	-0.017*** (0.006)	0.137*** (0.009)	0.149*** (0.009)
Overidentification test			0.000	0.000
1 st order serial correlation	0.000	0.000	0.000	0.000
2 nd order serial correlation	0.000	0.000	0.778	0.381
Observations	119,764	119,764	92,418	119,764
(b) Capital intensity				
Capital intensity (t-1)	0.971*** (0.001)	0.471*** (0.008)	0.757*** (0.035)	0.921*** (0.011)
Overidentification test			0.000	0.000
1 st order serial correlation	0.000	0.000	0.000	0.000
2 nd order serial correlation	0.000	0.000	0.476	0.692
Observations	119,764	119,764	93,513	119,764
(c) Market share				
Market share (t-1)	0.996*** (0.000)	0.486*** (0.007)	0.409*** (0.043)	0.935*** (0.011)
Overidentification test			0.000	0.000
1 st order serial correlation	0.000	0.000	0.000	0.000
2 nd order serial correlation	0.000	0.000	0.076	0.972
Observations	119,764	119,764	89,523	119,764

This table presents the results of AR(1) regressions. Column 1 presents estimates of ρ from the regression $y_{it} = \alpha + \rho y_{i,t-1} + \delta_t + v_{it}$ estimated using OLS, where y_{it} is the labour share in panel A, capital intensity in panel B, and market share in panel C. Column 2 presents within-group estimates of the model. Column 3 presents GMM estimates of the model in first differences, with lags of y_{it} dated $t - 2$ and earlier used as instruments for $\Delta y_{i,t-1}$. Column 4 presents system GMM estimates, where the lagged first difference of y_{it} is additionally used as an instrument for $y_{i,t-1}$ in the equation in levels. Serial correlation test: p-value for the null hypothesis of no serial correlation from the Arellano and Bond (1991) test for 1st and 2nd-order serial correlation. Overidentification test: p-value for the null hypothesis of instrument exogeneity from the Hansen test of overidentifying restrictions. Robust standard errors clustered by firm are in parentheses, with Windmeijer's (2005) finite-sample correction applied to GMM. * p<0.10, ** p<0.05, *** p<0.01

TABLE 2.A2: REDUCED-FORM EQUATIONS FOR 2012

	First Differences (1)	Levels (2)
(a) Labour share		
F-statistic	161.075	39.502
R ²	0.273	0.041
Observations	10,290	10,290
(b) Capital intensity		
F-statistic	16.244	268.467
R ²	0.021	0.138
Observations	10,942	10,942
(c) Market share		
F-statistic	3.059	41.253
R ²	0.011	0.025
Observations	8,417	8,417

This table presents reduced-form “first-stage” results for the last year of the sample when the greatest number of lagged values is available to serve as instruments. Column 1 (first differences) is an OLS regression of $\Delta x_{i,t-1}$ on $x_{i,t-2}, x_{i,t-3}, \dots, x_{i,t-7}$, where x is the labour share in panel A, capital intensity in panel B, and market share in panel C. Column 2 (levels) is an OLS regression of $x_{i,t-1}$ on $\Delta x_{i,t-2}, \Delta x_{i,t-3}, \dots, \Delta x_{i,t-6}$. The F-statistic reflects a test of the null hypothesis that the slope coefficients are jointly zero. * p<0.10, ** p<0.05, *** p<0.01

Appendix B: GMM assumptions and instrument validity

This Appendix discusses the conditions required for the consistency of the econometric estimates given in the paper and provides evidence that these conditions are satisfied. Consistency of the System GMM estimates relies on three assumptions. The first is instrument validity. Instruments are valid if they are strong predictors of the potentially endogenous regressors (relevance), while only affecting the current labour share through the endogenous variables (exclusion restriction). The second assumption is that the disturbances ε_{it} are serially uncorrelated. Third, the use of the levels equation requires an additional assumption of stationarity. I review each of these assumptions below.

I begin by testing whether the lags of the labour share, capital intensity, and market share satisfy the instrument relevance requirement. Consider the market share variable as an example. The market share at time $t - 2$ and $t - 3$ (instruments in the differenced equation) needs to be highly correlated with the year-over-year change in the market share observed at time t and $t - 1$ (the endogenous regressors). Similarly, in the levels equation, the lagged market share changes at time $t - 1$ and $t - 2$ need to be good predictors of the levels at time t and $t - 1$.

Table 2.B1 reports the results of first-stage regressions that correspond to the main GMM specification in column 5 of Table 2.3. Panel (a) considers the equation in differences, while panel (b) considers the equation in levels. Within each panel, the first row shows the F-statistic from the test of joint significance of the instruments in a regression of the endogenous variable on its lags without any controls. The second row reports the same test statistic after controlling year, sector, and region fixed effects. Finally, the third row controls additionally for all other instruments in the model. The third row represents a standard first stage equation in instrumental variable estimation.

Column (1) of panel (a) shows that the second and third lags of the labour share are strong predictors of the first lag of the change in the labour share. The F-statistic is 4,014.1 when all the controls and instruments are included. Columns (2) and (3) show the strength of the instruments for the current and lagged difference of capital intensity (F-statistics of 219.2 and 265.7, respectively). Finally, columns (4) and (5) show that the instruments for the current and lagged change in market share are also strong (F-statistics of 27.6 and 147.2, respectively). Similarly, in panel (b), which tests the relevance of instruments in the levels equation, all instruments are very strong predictors of the potentially endogenous regressors (F-statistics of 206.9 or greater). Therefore, the first-stage results indicate that inference in the GMM model is unlikely to be affected by problems caused by weak instruments.

The exclusion restriction is more difficult to test directly but can be assessed in two ways. The first is through the Hansen test of overidentifying restrictions. Table 2.B2 reports the results for the baseline specification. Part (a) shows that the null hypothesis of joint instrument validity cannot be rejected for the model as a whole (p-value 0.459). Parts (b) through (d) test several subsets of the instrument set for the levels equation and the first-difference equation separately. The results indicate that the satisfactory overall outcome of the Hansen test is not driven by any single component of the model.¹⁸

One potential concern with the Hansen test is that it can be weakened by the presence of many instruments. An alternative, intuitive way to evaluate the validity of the exclusion restriction consists of removing each candidate instrument from the GMM instrument matrix, including it directly in the estimated 'second-stage' model, and checking that it is insignificant, conditional on the remaining instruments. Table 2.B3 presents the results of this test. None of the coefficients on the tested instruments in columns (1) to (6)

¹⁸ In the case of capital intensity, the null hypothesis that the second and third lag are exogenous instruments for the difference equation is rejected at a significance level only slightly greater than 10%. Replacing the second lag with the fourth lag improves the Hansen test result for this subset of instruments (p-value 0.510) without materially affecting the estimated coefficients in the model or other diagnostic tests.

are significantly different from zero. This result does not suggest any correlation between the instruments and the error term of the estimated model, providing further reassurance about the validity of the instruments in the GMM model used to estimate the relationship between a firm's labour share and its capital intensity and market share.

Besides instrument validity, consistent estimation of the model relies on the assumption that the disturbances ε_{it} are serially uncorrelated, as specified by equation (2.6). This can be verified by testing for first and second-order serial correlation in the first-differenced residuals (Arellano and Bond, 1991). We would expect negative first-order serial correlation, but no second-order or higher-order serial correlation, in the first-differenced residuals, under the maintained assumptions. The results of these tests are reported in each of the regression tables discussed in Section 2.6. Together with the conclusions of this Appendix, they suggest that the moment conditions are well-specified.

The use of the levels equation requires two additional assumptions. First, the first differences of the two explanatory variables, capital intensity and market share, should be uncorrelated with the firm fixed effect α_i in equation (2.4). Second, the model specified by equation (2.4) needs to be the correct specification for labour shares “sufficiently long” before the first observation in the panel (Blundell and Bond, 2000).

Three tests suggest that this condition applies sufficiently well in the sample. First, limiting the sample to a balanced panel of firms that are present in every year of the sample period, and are more likely to be older and well-established, gives similar results as the main sample (see the discussion of robustness checks and Table 2.5 in Section 2.6). Second, the results are similar when limiting the sample directly to firms aged 10 years or more in the first year of the sample period, based on the year of incorporation (see Table 2.C1). Third, the validity of the instruments for the levels equation is tested specifically via a Difference Hansen test in Table 2.B2. None of these diagnostics highlight any concerns.

A final potential problem with GMM estimation relates to overfitting. As the time dimension of the panel increases, the set of possible instruments rises, since more past values of each variable become available. In practice, however, using too many instruments may bias the results (Roodman, 2009). While it is difficult to say exactly how many instruments are “too many,” symptoms of overfitting include an implausibly perfect p-value of 1.000 in the Hansen test of overidentifying restrictions. To avoid overfitting, I limit the number of instruments to the lagged levels dated $t - 2$ and $t - 3$ in the first-difference equation and lagged differences dated $t - 1$ and $t - 2$ in the levels equation, rather than using all past values. The first-stage results reported in Table 2.B1 and discussed in this Appendix show that these instruments are sufficiently strong to obviate the need for a greater number of them.

The empirical model, which flexibly incorporates the standard assumptions from the literature on estimating production functions, thus appears to be well-specified in responding to concerns about serial correlation and potential endogeneity of the variables of interest.

TABLE 2.B1: FIRST-STAGE REGRESSIONS

(a) Difference equation					
	Labour share	Capital intensity		Market share	
	$\Delta s_{i,t-1}^L$	$\Delta \left(\frac{K}{L}\right)_{it}$	$\Delta \left(\frac{K}{L}\right)_{i,t-1}$	Δs_{it}^M	$\Delta s_{i,t-1}^M$
	(1)	(2)	(3)	(4)	(5)
<i>Joint significance of the instruments: lagged levels (t-2 and t-3) (F-statistic):</i>					
(1) No Controls	4,447.1	150.5	181.1	19.9	88.8
(2) Year, Sector, Region FE	4,842.6	235.2	280.5	34.8	130.4
(3) All Instruments	4,014.1	219.2	265.7	27.6	147.2
(b) Levels equation					
	Labour share	Capital intensity		Market share	
	$s_{i,t-1}^L$	$\left(\frac{K}{L}\right)_{it}$	$\left(\frac{K}{L}\right)_{i,t-1}$	$s_{i,t-1}^M$	$s_{i,t-1}^M$
	(1)	(2)	(3)	(4)	(5)
<i>Joint significance of the instruments: lagged differences (t-1 and t-2) (F-statistic):</i>					
(1) No Controls	12,025.2	1,354.7	1,942.4	206.9	259.2
(2) Year, Sector, Region FE	12,990.9	1,321.9	2,007.9	362.6	477.6
(3) All Instruments	10,506.6	1,082.0	1,671.4	335.9	468.1

First-stage regressions for the endogenous variables in the main GMM specification in column 3 of Table 2.4. In panel (a), the dependent variables are in first differences and the instruments are the second and third lagged difference. In panel (b), the dependent variables are in levels and the instruments are the first and second lagged difference. F-statistics test the joint significance of the two instruments for each potentially endogenous regressor. Estimation method: least squares. Robust standard errors (in parentheses) are clustered by firm. The sample size is 119,764.

TABLE 2.B2: TESTS OF OVERIDENTIFYING RESTRICTIONS

	p-value (H ₀ : the instrument set is exogenous)
Hansen test	
(a) Overall model	0.459
Difference Hansen test for instrument subsets	
(b) Instruments for the difference equation	
(i) Labour share	0.275
(ii) Capital intensity	0.102
(iii) Market share	0.556
(c) Instruments for the levels equation	
(i) Labour share	0.392
(ii) Capital intensity	0.258
(iii) Market share	0.986
Observations	119,764

Part (a) reports the Hansen test of overidentifying restrictions for the main GMM specification in column 3 of Table 2.4. Parts (b) and (c) report the Difference Hansen tests of exogeneity for selected instrument subsets. Tests statistics are asymptotically chi-squared distributed with the degrees of freedom equal to the number of instruments in the given subset. The null hypothesis in each test is that the instruments are exogenous. The sample size is 119,764.


TABLE 2.B3: ADDITIONAL TESTS OF OVERIDENTIFICATION

(a) Difference equation						
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Instruments tested in the second stage: lagged levels (t-2 and t-3)</i>						
Labour share, $s_{i,t-2}^L$	0.002 (0.015)					
Labour share, $s_{i,t-3}^L$		0.025 (0.022)				
Capital intensity, $(K/L)_{i,t-2}$			0.005 (0.006)			
Capital intensity, $(K/L)_{i,t-3}$				0.003 (0.006)		
Market share, $s_{i,t-2}^M$					0.013 (0.012)	
Market share, $s_{i,t-3}^M$						-0.004 (0.020)
Serial correlation test						
1 st order	0.000	0.109	0.000	0.000	0.000	0.000
2 nd order	0.341	0.730	0.381	0.116	0.387	0.093
Overidentification test	0.818	1.000	0.841	0.915	0.635	0.828
(b) Levels equation						
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Instruments tested in the second stage: lagged differences (t-1 and t-2)</i>						
Labour share, $\Delta s_{i,t-1}^L$	-0.002 (0.015)					
Labour share, $\Delta s_{i,t-2}^L$		0.003 (0.008)				
Capital intensity, $\Delta(K/L)_{i,t-1}$			-0.005 (0.006)			
Capital intensity, $\Delta(K/L)_{i,t-2}$				0.001 (0.006)		
Market share, $\Delta s_{i,t-1}^M$					-0.013 (0.012)	
Market share, $\Delta s_{i,t-2}^M$						0.008 (0.011)
Serial correlation test						
1 st order	0.000	0.000	0.000	0.000	0.000	0.000
2 nd order	0.341	0.069	0.381	0.104	0.387	0.109
Overidentification test	0.818	0.844	0.841	0.753	0.635	0.311

This table presents additional, intuitive tests of instrument exogeneity. In each column one of the instruments from the main GMM specification in column 3 of Table 2.4 is excluded from the instrument matrix and entered directly into the estimated model. The serial correlation test is the Arellano and Bond (1991) test for first and second-order serial correlation (p-values reported). The overidentification test is the Hansen test of overidentifying restrictions (p-value reported). Standard errors clustered by firm are in parentheses, with Windmeijer's (2005) finite-sample correction applied. The sample size is 119,764. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix C: Additional figures and tables

FIGURE 2.C1: ALTERNATIVE MEASURES OF CAPITAL INTENSITY AND MARKET SHARE



Local polynomial plots of the log labour share against alternative measures of capital intensity and market share for the GMM estimation sample from column 3 of Table 2.4: (a) real intangible assets per employee, (b) real total fixed assets per employee, (c) market share in a 4-digit sector and NUTS3 region, (d) market share in a 4-digit sector and NUTS3 region for the sub-sample of firms in non-tradable sectors (52,564 obs.). The shaded areas correspond to 95% confidence intervals.

TABLE 2.C1: TESTING THE VALIDITY OF USING THE LEVELS EQUATION IN SYSTEM GMM

	Firms Aged 10+
	(1)
Capital Intensity	-0.067** (0.031)
Market Share	-0.213*** (0.036)
Serial correlation test	
1 st order	0.000
2 nd order	0.846
Overidentification test	0.339
Year FE	x
Sector FE	x
Region FE	x
Sample Size	83,485

This table tests the validity of using the levels equation in System GMM by checking whether the results are similar to the main specification in column 3 of Table 2.4 when the sample is limited to firms that have been incorporated for 10 years or more by the first year of the sample period. See the main text for details. Coefficients represent long-run elasticities. Serial correlation test: p-value for the null hypothesis of no serial correlation from the Arellano and Bond (1991) test for first and second-order serial correlation. Overidentification test: p-value for the null hypothesis of instrument exogeneity from the Hansen test of overidentifying restrictions. Robust standard errors clustered by firm are in parentheses, with Windmeijer's (2005) finite-sample correction applied to GMM. *** p<0.01, ** p<0.05, * p<0.10.

TABLE 2.C2: HIGH- AND LOW-WAGE FIRMS

Sector	Distribution of sample firms by sector		Data on the employee population by sector			
	Low-wage firms (%)	High-wage firms (%)	(a) Wages		(b) Skill level	
			Mean gross earnings (£/week)	Median gross earnings (£/week)	Employees with a college degree (%)	Average education (years)
A Agriculture, forestry, and fishing	2.1	1.0	324	288	14.7	12.5
B+D+E Mining, energy, and water supply	1.3	3.1	678	531	24.1	12.8
C Manufacturing	25.6	29.8	603	462	19.7	12.7
F Construction	8.7	13.4	579	500	17.2	12.4
G Wholesale, retail, repair of vehicles	30.8	22.9	298	243	13.7	12.5
H Transport and storage	3.6	3.2	495	423	13.4	12.3
I Accommodation and food services	12.3	0.8	212	157	14.1	12.9
J Information and communication	3.0	8.6	734	635	57.8	15.0
L Real estate activities	2.0	1.6	417	369	27.2	13.1
M Professional, scientific, technical activities	3.6	7.7	669	577	55.9	14.8
N Administrative and support services	7.1	8.0	373	308	18.8	12.8
Total	100.0	100.0				

This table shows the sectoral distributions of high- and low-wage firms in the sample and provides weighted population estimates of wages and qualifications by sector using the Labour Force Survey for the April-June quarter 2012. High-wage firms are defined as those whose total compensation per employee is greater than the year's median, on average during these firms' appearance in the sample. Years of education are calculated as the age when the individual's educational highest qualification was obtained minus five.

Appendix D: Mathematical derivations

(Not intended for publication)

This appendix contains the following supporting material:

- (a) derivations of equations (2.1), (2.3), and (2.5) in the main text;
- (b) calculations that support the interpretation of estimated elasticities in Section 2.6.

Mathematical derivations of key equations

Equation (2.1)

This equation relates the labour share of income to the partial elasticity of output with respect to labour and the price elasticity of demand in a simple, partial-equilibrium model of a profit-maximizing firm operating in an imperfectly competitive product market.

Consider firm i with the production function $Y_i = Y(A_i L_i, B_i K_i)$, where Y_i is output, L_i is labour, K_i is capital, and A_i and B_i are labour- and capital-augmenting productivity, respectively. Abstracting from the multiplicative productivity parameters for notational simplicity, the firm's profit maximization problem is

$$\max_{L_i, K_i} \pi_i = P(Y_i) Y(L_i, K_i) - w L_i - r K_i.$$

The first-order condition with respect to labour is

$$w = P'(Y_i) Y_L(L_i, K_i) Y(L_i, K_i) + P(Y_i) Y_L(L_i, K_i),$$

where $Y_L(L_i, K_i)$ is the partial derivative of output with respect to labour. Rearranging the right-hand side gives

$$\begin{aligned} w &= Y_L(L_i, K_i) [P(Y_i) + P'(Y_i) Y(L_i, K_i)] \\ &= Y_L(L_i, K_i) P(Y_i) \left(1 + \frac{P'(Y_i)}{P(Y_i)} Y(L_i, K_i) \right) \\ &= Y_L(L_i, K_i) P(Y_i) \left(1 + \frac{1}{\eta_i} \right) \end{aligned}$$

where $\eta_i = \frac{\partial Y_i}{\partial P_i} \frac{P_i}{Y_i}$ is the price elasticity of demand. In other words, the marginal cost of labour equals the marginal revenue product of labour times the mark-up. When this first-order condition is satisfied, the labour share is

$$\begin{aligned} s_i^L &\equiv \frac{wL_i}{P_i(Y_i)Y_i(L_i, K_i)} \\ &= \frac{L_i}{P_i(Y_i)Y_i(L_i, K_i)} \left(P_i(Y_i)Y_L(L_i, K_i) \left(1 + \frac{1}{\eta_i} \right) \right) \\ &= \frac{Y_L(L_i, K_i)L_i}{Y(L_i, K_i)} \left(1 + \frac{1}{\eta_i} \right) \end{aligned}$$

Setting $\varepsilon_i^{YL} = Y_L(L_i, K_i) \frac{L_i}{Y_i}$, the partial elasticity of output with respect to labour, the labour share is thus given by

$$s_i^L = \varepsilon_i^{YL} \left(1 + \frac{1}{\eta_i} \right) \quad (2.1)$$

as shown in the main text.

If the production function is Cobb-Douglas, then ε_i^{YL} is constant, and the labour share will only be a function of the price elasticity of demand. If the production function is not Cobb-Douglas, then ε_i^{YL} will be a function of capital, labour, and the productivity parameters, such that the labour share will depend on both the price elasticity of demand and the factor inputs.

The derivation of equation (2.8) in the main text follows a similar process, allowing additionally for imperfect competition in the labour market.

Equation (2.3)

This equation relates a firm's market share to the price elasticity of demand in an illustrative version of a Cournot model with n profit-maximizing firms, each producing a

homogeneous product with marginal cost c_i , and linear inverse demand, $P = a - bQ$, where $Q = \sum_i q_i = q_i + q_{-i}$ is total quantity produced.

The market share of firm i is the ratio of the quantity produced by that firm to total quantity produced in the market,

$$s_i^M \equiv \frac{q_i}{Q}$$

The Nash equilibrium solution for each firm's output q_i and total quantity Q arises from the optimal response of each firm to the choices of the remaining firms. In this model, firm i 's profit is

$$\pi_i = P(q_i + q_{-i})q_i - c_i q_i.$$

The profit-maximizing first-order condition is thus

$$\frac{\partial P}{\partial q_i} q_i + P - c_i = 0.$$

Given linear demand, this is equivalent to

$$-bq_i + (a - b(q_i + q_{-i})) - c_i = 0. \quad (2.D.1)$$

Thus, as a function of total quantity produced, and using $Q = q_i + q_{-i}$, firm i produces

$$q_i = \frac{a - c_i - bQ}{b}. \quad (2.D.2)$$

To solve for the total quantity produced in the market as a function of the model parameters, note that equation (2.D.1) yields the system

$$-bq_1 + (a - bQ) - c_1 = 0$$

$$-bq_2 + (a - bQ) - c_2 = 0$$

$$\vdots$$

$$-bq_n + (a - bQ) - c_n = 0$$

Adding these n equations together pins down the relationship between total quantity produced, the number of firms, marginal costs, and the slope and intercept of the demand curve:

$$-bQ + n(a - bQ) - n\bar{c} = 0,$$

where $\bar{c} = \frac{1}{n} \sum_{i=1}^n c_i$ is average marginal cost in the market. Rearranging for Q , total quantity produced is thus

$$Q = \frac{n}{n+1} \frac{a - \bar{c}}{b}. \quad (2.D.3)$$

Combining (2.D.2) and (2.D.3) gives an expression for firm i 's market share as

$$s_i^M \equiv \frac{q_i}{\sum_i q_i} = \frac{n+1}{n} \frac{a - c_i}{a - \bar{c}} - 1. \quad (2.D.4)$$

To see how the market share can be expressed in terms of the price elasticity of demand η , note that

$$\begin{aligned} \eta &\equiv \frac{dQ}{dP} \frac{P}{Q} \\ &= -\frac{1}{b} \left(\frac{a - bQ}{Q} \right) \\ &= 1 - \frac{a}{bQ} \\ &= 1 - \frac{a}{\frac{n}{n+1} (a - \bar{c})} \\ &= -\frac{a + n\bar{c}}{n(a - \bar{c})} \end{aligned}$$

Rearranging,

$$n(a - \bar{c}) = -\frac{a + n\bar{c}}{\eta}.$$

Substituting into (2.D.4), the market share can be expressed as

$$s_i^M = \frac{(n+1)(a-c_i)}{a+n\bar{c}}(-\eta) - 1. \quad (2.3)$$

This illustrates a possible link between firm market shares and the price elasticity of demand in an instance of a model of imperfect competition.

Equation (2.5)

The autoregressive distributed lag model given by equation (2.5) can be derived by imposing a specific form of serial correlation on the time-varying component of the error term, v_{it} , in equation (2.4). Suppose that this variable follows an AR(1) process, such that we have a model given by:

$$s_{it}^L = \beta_1 \left(\frac{K_{it}}{L_{it}} \right) + \beta_2 s_{it}^M + \theta_j + \varphi_r + \delta_t + (a_i + v_{it}) \quad (2.4)$$

$$v_{it} = \rho v_{i,t-1} + \varepsilon_{it} \quad (2.4b)$$

$$\varepsilon_{it} = MA(0) \quad (2.4c)$$

Solving equation (2.4) for v_{it} and solving its lagged version for $v_{i,t-1}$ gives

$$v_{it} = s_{it}^L - \beta_1 \left(\frac{K_{it}}{L_{it}} \right) - \beta_2 s_{it}^M - \theta_j - \varphi_r - \delta_t - a_i$$

and

$$v_{i,t-1} = s_{i,t-1}^L - \beta_1 \left(\frac{K_{i,t-1}}{L_{i,t-1}} \right) - \beta_2 s_{i,t-1}^M - \theta_j - \varphi_r - \delta_{t-1} - a_i.$$

Substituting these expressions into equation (2.4b) gives the autoregressive distributed lag specification,

$$\begin{aligned} s_{it}^L &= \rho s_{i,t-1}^L + \beta_1 \left(\frac{K_{it}}{L_{it}} \right) - \rho \beta_1 \left(\frac{K_{i,t-1}}{L_{i,t-1}} \right) + \beta_2 s_{it}^M - \rho \beta_2 s_{i,t-1}^M + (\delta_t - \rho \delta_{t-1}) \\ &\quad + (1 - \rho)(\theta_j + \varphi_r + a_i) + \varepsilon_{it}. \end{aligned}$$

First-differencing to remove the time-invariant firm effect, we obtain

$$\Delta s_{it}^L = \rho \Delta s_{i,t-1}^L + \beta_1 \Delta \left(\frac{K_{it}}{L_{it}} \right) - \rho \beta_1 \Delta \left(\frac{K_{i,t-1}}{L_{i,t-1}} \right) + \beta_2 \Delta s_{it}^M - \rho \beta_2 \Delta s_{i,t-1}^M + \Delta \varepsilon_{it}.$$

This is equivalent to

$$\Delta s_{it}^L = \rho \Delta s_{i,t-1}^L + \pi_1 \Delta \left(\frac{K_{it}}{L_{it}} \right) + \pi_2 \Delta \left(\frac{K_{i,t-1}}{L_{i,t-1}} \right) + \pi_3 \Delta s_{it}^M + \pi_4 \Delta s_{i,t-1}^M + \Delta \varepsilon_{it}, \quad (2.7)$$

where $\pi_1 = \beta_1$, $\pi_2 = -\rho \beta_1$, $\pi_3 = \beta_2$, and $\pi_4 = -\rho \beta_2$. This equation can be estimated consistently using GMM, as described in the main text.

Interpreting the estimated elasticities in Section 2.6

Section 2.6 interprets the estimated long-term elasticities from Column 3 of Table 2.7 by (a) calculating the estimated labour share reduction that would result from a 1 percentage-point increase in the market share of an average firm in the sample, and (b) calculating the estimated increase in net tangible assets per employee that would be required to reduce the labour share by the same amount as in (a).

The corresponding calculations are as follows. The mean market share in the 4-digit sector is 1.6%. Hence, a 1.2 percentage point rise in the market share corresponds to a 74% increase for the average firm. Multiplying the estimated elasticity of -0.161 by 74% implies a 12% reduction in the labour share. Since the mean labour share in the data is 0.83, this suggests a 10 percentage-point decline in the labour share of the average firm, holding capital intensity constant.

Similarly, a £42,000 increase in net tangible assets per employee is a 123% increase for the average low-wage firm, given that the mean of this variable in the sample of low-wage firms is £34,063. Multiplying the estimated elasticity of -0.096 by 123% gives a 12% reduction in the labour share, which, as above, corresponds to 10 percentage points for the average low-wage firm with the labour share of 0.84. This increase in net tangible assets is approximately 1.4 times the median annual employee compensation cost in the sample of

firms (£29,508). The calculations for the average high-wage firm are carried out in a similar way, based on mean net tangible assets per employee of £44,500 and mean labour share of 0.82.

Section 2.6 also calculates the estimated percentage-point reduction in the labour share for a 1 standard-deviation increase in market share or capital intensity for the average firm. In the sample, the mean market share is 1.6% and the standard deviation is 4.4%. Dividing the standard deviation by the mean and multiplying by the estimated elasticity of -0.161 gives a 44% reduction in the labour share. For the average firm with the labour share of 0.83, this implies a reduction of 36.8 percentage points.

Similarly, the standard deviation of capital intensity (£81,292 and £117,114 for low-wage and high-wage firms, respectively) divided by the mean (£34,063 and £44,500) and multiplied by the estimated elasticities (-0.096 and -0.054) gives a reduction in the labour share of 23% and 14%, respectively. This corresponds to a 19.3 and 11.7 percentage-point fall in the mean labour share for the average low-wage and high-wage firm in the sample, respectively.

Chapter 3

Pension Shocks and Wages*

Co-authored with Brian Bell
(King's Business School and Centre for Economic Performance,
London School of Economics)

Abstract

How do wages respond to firm-level idiosyncratic cost shocks? We create a unique dataset that links longitudinal data on workers' compensation to the unexpected costs that UK firms have been forced to pay to plug large deficits in their legacy defined benefit pension plans. We show that firms are able to share the burden of such costs when a significant share of their workers are current or former members of the plan. We also investigate how compensation responds to the closure of defined benefit plans to future benefit accrual. We find that firms are able to use such closures to effectively reduce total compensation of workers who are plan members. These results point to significant frictions in the labour market, which we show are a direct result of the pension arrangement that workers have. Closing schemes has an implicit cost for firms since it reduces the frictions that workers face.

Keywords: Wages, Pensions, Frictions

JEL Codes: J31, J32, G32

* We are grateful to Steve Bond, Stephen Machin, and seminar participants at Birmingham, Cardiff, Manchester, Sheffield, Surrey and UCL/LSE. This paper contains data from the UK Office for National Statistics (ONS) supplied by the UK Data Service (UKDS), which are Crown copyright and are reproduced with the permission of the Controller of Her Majesty's Stationery Office and the Queen's Printer for Scotland. This work uses research datasets that may not exactly reproduce National Statistics aggregates. ONS and UKDS bear no responsibility in relation to the interpretation or further analysis of the data. All results have been reviewed to ensure that no confidential information is disclosed. This research was funded by the ESRC and the John Fell Fund.

3.1. Introduction

There is a growing body of academic work that examines the impact of firm-specific shocks on wages. In almost all cases, this evidence relates to shocks to demand or profitability that are thought to be plausibly exogenous for wages. There is, for example, a large literature in labour economics that has focused on the link between firm profitability and wages, often understood within a rent-sharing context. Van Reenen (1996) shows that firms that innovate pay higher wages, with innovation being viewed as a good instrument for rents. In a similar vein, Kline, Petkova, Williams and Zidar (2017) show that firms that generate ex-ante valuable patents – relative to a control group of rejected patent applicant firms – share some of the surplus with workers. On average, workers capture 29 cents of every dollar of patent-induced surplus. Using a different identification approach, Garin and Silverio (2017) examine how sensitive wages are to idiosyncratic export demand shocks using Portuguese matched worker-firm data. They find that a shock that reduces a firm's value-added by 10 percent reduces the wages of incumbent workers by 1.5 percent. They additionally show that these effects are stronger in industries with longer worker tenure and lower turnover rates, suggesting that mobility frictions may be an important explanation for these findings.

This paper adds to the literature by focusing instead on a firm-specific cost shock – specifically, a shock to non-wage employment costs. On average, the non-wage elements of compensation account for around 20% of total worker remuneration.¹ The components of these additional costs differ across countries and over time, but key elements include employer pension contributions, private healthcare and employment taxes. The focus of this paper is on how unexpected shocks to the level of one of those non-wage components of labour cost are borne by the firm and the workers. Importantly, this cost shock differs

¹ In 2016, non-wage labour costs represented 17.3% of total compensation in the UK, 18.9% in the US, and 22.2% in the Eurozone (OECD, 2017).

dramatically across firms and is independent of current employment and wages as it arises from historical commitments that firms made.

The context for our analysis is the dramatic shift in private pension provision and the pension-related costs of UK firms over the last two decades. Pensions can essentially be divided into two types: defined benefit (DB) and defined contribution (DC). DB schemes provide members with a guaranteed lifetime income in retirement, typically a percentage of their final or average salary. DC schemes, on the other hand, are like individual savings accounts, where the size of the pension pot at the time of retirement depends on individual contributions and accumulated returns. Historically, UK firms that provided pensions as part of the remuneration package offered DB schemes. At the start of the 1980s, 97% of workers in the private sector that were active members of an employer pension scheme were in DB schemes. Whilst this figure declined as DC schemes became more commonly available, it was still 81% in 2000. Yet by 2015, it had fallen to 29%. Figure 3.1 shows these abrupt shifts in terms of the total number of workers. Interestingly, this is an entirely private-sector development – public sector workers have continued to have access to DB schemes, though often the generosity of these schemes has changed. A similarly dramatic move away from DB schemes occurred in the United States (Haverstick, Munnell, Sanzenbacher and Soto, 2010).

What has caused this precipitous fall in DB membership in the private sector since the start of the millennium? The key to understanding this change is to focus on the assets and liabilities of DB schemes. On the asset side, firms (and workers) make ongoing contributions to the scheme that are then invested in a portfolio of assets – typically a mix of equities, bonds and property. The value of the assets thus crucially depends on asset returns. To the extent that such returns are poor, there must be additional contributions from firms and/or workers to achieve a given asset level. On the liabilities side, there are two

key elements. First, the longevity of workers affects the cost of the scheme, since workers are guaranteed a stream of income until death. Second, the discount rate affects the calculation of the present value of these liabilities. As the discount rate falls, the present value of liabilities rises – sometimes dramatically.

The last two decades have witnessed a perfect storm in all the elements that make up these calculations. Since the abolishment of dividend tax credits for UK pension schemes in 1997, rising longevity, low interest rates that fell further following the 2008 financial crisis, and long periods of weak equity returns have depressed pension scheme assets and inflated the liabilities. As these market risks tend to be largely unhedged at the pension scheme level, they have led to a steep increase in deficits and, as a result, firms have begun to close their DB schemes. By 2016, 85% of DB schemes were closed to new members and 35% were also closed to future accrual of benefits.

What can firms do under these circumstances? An important point to note is that, as a result of pension fund legislation, firms are unable to renege on the DB commitments historically made to workers, nor are they able to decide for themselves how to deal with the deficit black holes that have developed in the schemes. The rules relating to DB pension plans in the UK are currently governed by the 2004 Pensions Act. A key aim of this Act is to ensure that DB schemes are funded appropriately given the expected value of liabilities and reasonable assumptions regarding asset returns. DB schemes are required to have a formal actuarial valuation every three years. This formal valuation determines the funding level of the scheme. The valuation triggers increased contributions by the sponsoring employer if a deficit exists and is unlikely to be eliminated without action. The trustees of the pension scheme must negotiate with the employer to establish a recovery plan. The resulting plan generally involves a sequence of deficit payments over subsequent years, which must satisfy the independent Pension Regulator.

The relative importance of deficit payments can be seen in Figure 3.2. During the 1990s, the majority of schemes were in surplus and thus required no such payments. On average, only 17% of total contributions were for deficit recovery. Since the start of the 2000s, this has risen to 32% and reached a peak of 42% in 2011. An alternative calculation simply asks how much in total has been spent on deficit contributions since 2000. The data underlying Figure 3.2 show that this equals £167bn. This steep rise in pension contributions is the main driver behind the sharp increase in the proportion of non-wage labour costs in total compensation that we observe in Figure 3.3 for the UK since 2000. Whilst these substantial deficit payments must be made by the firm, it is unclear who bears the ultimate burden: the firm, the workers, or both.

Firms can, however, change the future commitments they make, by adopting one of two approaches. First, they can close the DB scheme to new workers – potentially offering them membership of a DC scheme as an alternative. However, workers already in the DB scheme would continue to accrue benefits under the scheme. Second, firms can close the DB scheme completely – to both new and existing workers. All workers could be offered a DC scheme as an alternative and no future accrual of benefits would occur. Although closing the scheme lets the firm avoid making any new DB pension commitments, it still has to honour the commitments already made to current and past workers – it merely avoids making any new ones. There are no legal restrictions on firms following either of these strategies. In practice, most firms have adopted a two-stage approach. In the early 2000s they began to close the DB schemes to new members. As the deficits continued to rise, firms responded by closing the schemes completely to future benefit accruals.²

² Some firms have also attempted to cap the liabilities related to their past commitments by offering members a cash lump sum to quit the scheme entirely. Others have insured themselves, partly or wholly, against future cost increases by purchasing third-party insurance. These actions are expensive, as the discount rates applied to the liabilities are typically higher than those applied for the purpose of calculating ongoing funding status.

The analysis in this paper therefore focuses on two aspects of pension costs. First, we explore the impact on wages of the deficit payments that firms are required to make to plug the black holes in the schemes. Do current workers bear any of this burden and if they do, what share do they bear? Second, we examine the impact of DB scheme closure on workers' remuneration. In a standard frictionless competitive model, in which workers and firms value pension contributions in the same way, the elimination of employer DB contributions should be matched by either an equivalent rise in DC contributions or a rise in workers' wages. By contrast, a labour market characterized by rent-sharing or by frictions—caused, for example, by firm-specific human capital—will generally not give rise to exactly offsetting changes in other components of compensation. We investigate both the sharing of deficit costs and the impact of scheme closures by constructing a unique dataset that links longitudinal data on workers' compensation to data on the DB pension schemes of UK firms.

The literature provides limited evidence on how individual wages respond to shocks to pensions. Rauh, Stefanescu and Zeldes (2017) examine the second of our two impacts for US firms by focusing on the effect of the freezing of pension benefit accruals on total payroll costs. The authors find that firms are able to achieve net cost savings by closing their DB schemes, suggesting that employees are not fully compensated for the loss of future benefits through higher wages or higher contributions to DC schemes. Whilst we find similar effects for the UK, our paper differs in two important ways. First, we use worker-level data to estimate the effects of pension benefit changes on wages. This allows us to account for personal characteristics and individual components of compensation, rather than rely on the rather crude measure of average firm-level wages available from company accounts. Second, we are able to explore the reasons for this result by providing evidence of substantial frictions in worker mobility generated by the pension arrangements.

The closure of schemes is shown to alter these frictions significantly and highlights a potential cost of scheme closure for the firm.

Our paper also complements the limited literature that measures the effect of DB pension contributions on other firm-level outcomes such as investment expenditure and dividends (Rauh, 2006; Phan and Hegde, 2012; Liu and Tonks, 2013; Chaudhry, Au Yong and Veld, 2017; Bunn, Mizen and Smietanka, 2018). We differ in our approach by focusing on wages and combining firm and individual-level data. Our analysis also has similarities with the work of Gruber (1997), whose analysis focused on the wage response to a large exogenous reduction in the payroll tax in Chile. Though the reduction was common across firms, the effective rate that applied differed across firms depending on the type of workers employed. While the legal incidence of the tax fell on employers, Gruber shows that the entire gain from the reduced payroll tax accrued to workers through higher wages. In our setting, the deficit payments required to support DB schemes are legally required to be paid by the firm. However, there is nothing to prevent this cost from being shared ex-post by workers. That is what we investigate in this paper.

The remainder of the paper is structured as follows. In Section 3.2, we describe the data sources and empirical strategy. Section 3.3 presents evidence on the impact of the cost shocks on wages, whilst Section 3.4 investigates how total compensation responds when firms close down pension schemes. In Section 3.5 we examine the source of the potential labour market frictions that our results imply. Section 3.6 concludes.

3.2. Data and empirical strategy

Data

We have collected data on pension schemes and payments from the annual reports of 475 UK-listed companies. The sampling frame requires companies to have been amongst the

300 largest UK-domiciled firms ranked by year-end market capitalization at any point over the period from 2000 to 2010. Of the 475 firms, 65% report exposure to at least one UK DB scheme – with the remaining 35% having either no pension exposure or only DC schemes. For each firm with a DB scheme, we attempt to collect data from the annual reports on: (i) total employer contributions, (ii) current service costs (defined below), (iii) year-end assets, liabilities and surplus/deficit, (iv) the triennial valuations and (v) scheme closure dates if any. This data has been consistently reported since 2001 when accounting rules were changed to require more detailed reporting of pension exposure. We collect data only on the UK pension exposure rather than the exposure to all schemes, as many firms have DB schemes in other countries. Figure 3.4 shows that the trends in total DB contributions in our sample follow a similar pattern to those for UK firms as a whole – and account for around one-third of total contributions.

Our wage data for workers comes from the Annual Survey of Hours and Earnings (ASHE). These data are a panel of 1% of employees, sampled randomly based on their unique social security number. The survey is conducted every April and the data are collected directly from employers' payrolls, so they are both reliable and benefit from a high response rate – response is, in fact, legally required. For our purposes the strength of the data is that firms report both gross wages and, since 2005, employer and employee pension contributions – as well as the type of pension scheme paid into. We use data from 2002 to 2016, as 2002 was the first year that a firm identifier was included for each worker, allowing us to match individual wages to the firm-level pension data.

To match the wage data for workers to our listed firm sample, we use the Dun and Bradstreet code. Of the 475 listed-firms, we find at least one worker in ASHE for 393 firms. There are two key reasons why we do not match workers for every firm. First, since ASHE is only a 1% sample, firms with small employment levels will frequently not have an

employee with the relevant social security number. Second, and more importantly, some firms listed and domiciled in the UK have almost their entire operation outside of the UK. This is particularly true for energy companies. Since ASHE only covers UK workers, such firms' employees will not be in the data. Finally, some firms do not disclose sufficient data about their UK defined-benefit pension schemes. This leaves us with a sample of 330 firms. Table 3.1 provides some summary statistics on the final sample of firms and workers.

We need a measure of the deficit payments to capture the firm-level cost shock and information on the closure of the pension scheme. The latter is easily identified as most firms report the closure date directly in the annual report. The only difficulty is if firms have multiple UK schemes and close them at different dates. The ASHE data does not allow us to distinguish which particular DB scheme the worker is enrolled in, so we cannot then precisely link the closure of a particular scheme to the individual worker. We take the conservative approach of using the closure date for the final UK scheme if more than one exists.

Firms are not required to report regular and deficit contributions separately, and when they do, the decision on how to classify contributions into the two types is not governed by a specific accounting standard. We therefore define the deficit payments that firms are required to make to their DB schemes as:

$$\text{Deficit Payment} = \text{Total Employer Contributions} - \text{Current Service Cost}$$

Current Service Cost is defined as “the increase in the present value of a defined benefit obligation resulting from employee service in the current period” (IAS 19). It therefore represents the current actuarial estimate of the cost of providing a DB scheme for the financial year for the current employees. It excludes the cost of any re-evaluation of the present value of the obligations for previous employees (or previous years of service for current employees). If the scheme has been closed to future accrual, the current service cost

is zero. To assess the accuracy of this approach, Figure 3.5 shows that for the 111 firms in our dataset that do voluntarily break out deficit contributions in their annual report, our measure of deficit payments matches the firms' own classification very closely. Like in firms' own reports, we estimate relatively stable aggregate deficit payments during the early 2000s, followed by a steep rise during the financial crisis and an improvement from 2013 onwards as more firms close their DB schemes. These trends are similar to aggregate estimates from the national accounts, shown earlier in Figure 3.2.

Appendix A provides additional details on the sample, including a breakdown of observation by year and workers' pension membership status, as well as a description of the key variables used in the empirical analysis.

Empirical strategy

Our empirical approach exploits the panel nature of the data to identify the effect of DB deficit cost shocks using within-firm variation over time. We estimate models of the form:

$$y_{ijt} = \alpha_i + \beta_j + \gamma_t + \sum_{k=1}^2 \delta_k DB_Deficit_Payment_{jt-k} + \pi X_{ijt} + \varepsilon_{ijt} \quad (3.1)$$

where y_{ijt} is a measure of remuneration for individual i , in firm j , at time t . We control for individual- and firm-fixed effects (respectively α_i and β_j), and further report estimates with match fixed effects. We also allow for common time shocks, γ_t , and a set of other observables, X_{ijt} that control for time-varying individual and firm characteristics. Our parameters of interest, δ_k , measure the effect on the outcome variable of up to two lags of the DB deficit payment measure (Table 3.B1 in the Appendix shows that the third lag does not have an economically or statistically significant effect on wages). To generate a measure that is both comparable across firms and easily interpretable, we deflate the DB deficit payment measure by the initial level of employment for the first observation of the firm in our sample. It is therefore a measure of the annual deficit payment per worker, and

the distribution of this variable is shown in Figure 3.6. It is clear that there is substantial variation in deficit payments per workers across firm-year observations. Table 3.B2 in the Appendix shows that the main results of the paper hold when using the first lag of employment as an alternative way to deflate DB deficit payments.

To assess the effects of scheme closure, we add an indicator variable, ω_{jt} , which equals 1 after the firm's scheme is closed to future accruals, to equation (3.1). We use either log wages, employer contributions, or employee contributions as the dependent variable, y_{ijt} . This enables us to measure how the individual components of employee compensation are affected by the firm's decision to close the scheme.

One potential concern is that the firm's decision to close the scheme may be jointly determined with wages, for instance in response to an external shock. To examine this potential endogeneity, we exploit the panel and allow for the indicator variable to affect y_{ijt} both before and after the actual closure date. Letting $t = 0$ be the actual closure date, we estimate:

$$y_{ijt} = \alpha_i + \beta_j + \gamma_t + \sum_{k=-4}^4 \theta_k \omega_{jt+k} + \pi X_{ijt} + \varepsilon_{ijt} \quad (3.2)$$

When $k < 0$, the estimated coefficients θ_k measure the impact of future scheme closure on wages or pension contributions and therefore provide evidence on whether elements of compensation have a pre-trend for those firms that close their scheme relative to a control group of firms that do not. For $k \geq 0$, the estimated θ_k measure the impact of actual scheme closure on the components of remuneration. By estimating separate coefficients for each year since closure, we can test whether any effect is immediate and sustained or whether effects grow over time. Rauh et al (2017) account for the potential endogeneity of scheme closure by using a propensity score matching approach to find “equivalent” firms that do not close their DB scheme at the same time. The approach we adopt here is essentially the

same since the extensive set of controls in the regression, e.g., time interacted with industry dummies, means that we are comparing treatment and control firms that are similar.

To further address the potential endogeneity of ω_{jt} , we assess whether the impact of DB scheme closure on pension contributions and wages differs for firms at which scheme closure coincides with the triennial actuarial valuation, compared to firms that close their DB schemes at other times. The actuarial valuation happens every three years based on a pre-determined schedule and is the basis for scheme trustees to require additional contributions from the firm. Firms that close their DB scheme in the same year as the valuation are therefore more likely to be doing so in response to exogenous changes in the funding level, rather than as a result of other considerations. This hypothesis is consistent with the cross-sectional findings of Munnell and Soto (2007) from the US, who find that the firm's decision to close the DB scheme is correlated with the underfunding level.

3.3. Firm-level cost shocks and wages

We begin our results by examining the effect of the deficits on the hourly wages of workers. We estimate fixed-effect panel wage regressions as in (1), with $\ln(\text{hourly wages})$ as the dependent variable. The results are presented in Table 3.2. As we move across the columns 1 to 4, we include increasingly more controls to test the robustness of the results. Column 5 then uses $\ln(\text{weekly earnings})$ as an alternative measure of compensation available in ASHE that takes into account the number of hours worked. Comparing the results across the columns, it is clear that the estimates are similar regardless of which control variables or specifications are used. This supports the contention that the deficit payments are indeed firm-specific shocks that are exogenous to other factors that affect wages. Overall, we find a statistically significant reduction in wages as a result of the deficit payments, suggesting that firms share some of the burden with workers.

We might, however, expect the impact of deficit payments on workers to depend on the extent to which the firm's labour force is exposed to the DB pension scheme. If the legacy costs that are being incurred are partly a result of historic promises made to current workers, it may prove easier for the firm to shift some of the cost burden than if the costs are a result of promises made to workers who are no longer with the firm. This is examined in Table 3.3. We take two approaches. First, we begin by estimating how the impact differs depending on whether the workers are themselves members of the DB scheme. To do so, we identify two mutually-exclusive groups of workers. The first group are DB members, defined as workers who are currently members of the scheme (active members) or have been members in the past while working at the same firm (deferred members). The second group are workers who have never belonged to the firm's DB scheme. The second approach is to classify firms into those that have large active membership (defined as the proportion of total pension scheme members who are still accruing benefits in the firm being above the median) and firms that have small active membership (i.e. most of the scheme members have retired or are deferred members). The reason for using this approach is that in practice, firms may not be able to set wages differentially for individuals based on their pension membership. Columns 1 to 3 report the results by individual worker pension membership status, while columns 4 to 6 report the results by active membership at the firm level. Each row in the table reports the sum of the coefficients (and associated standard error) on two lags of the DB deficit measure for each type of worker. Again, as we go across the table we include more controls to test the robustness of the results.

Column (3) reports the impact on scheme members controlling for match fixed-effects, measures of firm performance, and industry*year dummies. The two measures of firm performance are two lags of the log of firm sales, which capture growth, and two lags of profitability, defined as gross value added as a percentage of sales. With these controls

in place, we find a strongly significant negative effect of the deficit payments on the wages of workers who are DB scheme members. In contrast, there is no impact on non-members. To assess the size of the estimated effect, Table 3.3 also presents estimates of the “sharing rate.” This is calculated for the current members using the mean level of annual deficit payments per worker in the sample (£1,561) and calculated using mean wages and hours. It is therefore an estimate of what proportion of the deficit payment is accounted for by reduced wages for these workers. We estimate a reduction in wages of around £144 per year, implying that these workers pay for about 10% of the deficit costs. The results are similar when we consider the wages of all workers at firms with large active membership. These sharing estimates are on the low-side of those reported using demand shocks which suggest a range of 15-30%. One important factor that differs between these studies is that the cost shock here does not relate to the current activity of the firm but to historic liabilities, whilst shocks to exports or innovation relate more directly to the efforts of the current workforce.

Do these magnitudes seem sensible? One way of addressing this is to compare them with the share of costs that workers in DB schemes commonly pay in accruals. In general, our data show that normal worker contributions to DB schemes account for around 30% of the total normal contributions, i.e., current members pay for around one-third of the normal costs of a DB scheme. So, an estimate of a sharing rate of around 10% for the deficit payments suggests that workers are paying a much lower share of those costs. This is unsurprising, given that many of the beneficiaries of the schemes are no longer with the firm, and in any case the accrued benefits for current workers are protected by law.

3.4. Impact of scheme closure

We now consider the second type of pension shock, namely the closure of the scheme to future accrual and its impact on the two relevant components of employee remuneration: employer pension contributions and wages. In Table 3.4, we present estimates of the employer pension contribution rate (as a percentage of gross pay) in which we include a dummy variable that turns on when the DB scheme closes. It therefore measures the long-run change in contribution rates from closing the DB scheme. Note that after the closure the rate can either fall to zero (if the employer decides to no longer offer any pension provision—this occurs in fewer than 1% of cases in the data) or to any rate that the employer chooses for a DC scheme.

We distinguish between workers that are exposed to the closure of the scheme and all other workers. In column 1, we define exposed workers simply as members of the DB scheme. In column 2, we define exposed workers as DB scheme members whose employer reported a non-zero DB pension contribution at some point in the past, while in column 3, we restrict this definition further to workers whose employer reported a positive contribution in the year the scheme closed. Overall, we find that employer contributions fall by 6-7 percentage points for scheme members, whilst those not currently exposed to the scheme see no significant change in contribution rates. The last column restricts the group of exposed workers even further to cases where the scheme closed in the year corresponding to the triennial valuation. This is to address potential concerns about the endogeneity of the decision to close the scheme. Here we see a fall in contributions of 14 percentage points as well as a marginally significant negative impact of 3 percentage points on all other workers.

Table 3.5 reports identical specifications except that we now have log hourly wages as the dependent variable. Given the results in Table 3.4, we might expect a significant

increase in wages for DB members when the scheme is closed to offset the decline in compensation coming from the 6-7% reduction in employer pension contributions³. However we can see that regardless of the exact specification used, there is no evidence of any offset in wages – none of the coefficients are statistically different from zero and they are all small in economic magnitude. So DB workers see a substantial fall in total compensation when the firm closes down the DB scheme.

Figures 3.7 and 3.8 illustrate how the closure of the DB scheme leads to a reduction in employer contributions with no offsetting increase in wages, and show that there is no evidence that this is because of endogeneity of scheme closure. These figures plot employer contributions and wages, respectively, for exposed DB members and all other workers for four years prior to and subsequent to scheme closure. Exposed DB members are defined as those who received a positive contribution from the firm in the year when the scheme closed. The figures show no evidence of trends in pension contributions or wages in the years preceding the closure of the scheme or any differences in remuneration between exposed and non-exposed workers, supporting the interpretation of the results in Tables 3.4 and 3.5. Detailed results supporting these plots are shown in Table 3.B3 in the Appendix. We have also examined whether there are different pre-trends in other indicators of firm performance, such as profitability, sales, total employment, for treated and control firms and find no evidence to support such a concern.

3.5. Frictions and worker mobility

Our results thus far point to significant costs being borne by workers who work in firms affected by pension deficits and scheme closures. These costs generally fall on those

³ We also estimated models with employee contributions as the dependent variable but these were all insignificant – they are available from the authors upon request.

exposed to the DB scheme and this suggests that frictions for these workers give firms some monopsony power. An obvious source of these frictions is the DB scheme itself. Workers in final salary DB schemes have an incentive to remain with a firm that continues to allow accruals because the pension entitlement generally depends upon the final salary of the worker which will grow both because of inflation and real wage growth. If the worker leaves the firm today, the final salary will be calculated by taking the current salary and uprating to account for inflation only, so any real wage growth effect will be lost. To give a concrete example, suppose a worker aged 40 has accrued 15 years of DB entitlement and will retire at age 65. Her current salary is £50,000 per year and she does not expect future promotions. The DB scheme is final-salary, accrues at the rate of $1/80^{\text{th}}$ per year, inflation is 2% and real wage growth in the economy is 1%. Ignoring future accruals, if the worker leaves today, she will receive a pension of £15,380 per year from age 65 (which is £9,375 in today's prices). If instead she remains, the pension would be £19,629 (£12,023 in today's prices) – and this of course ignores the 25 years of future accrual. This illustrates the increased mobility costs for DB workers compared to otherwise identical workers in a DC scheme, whose pension pay-out is wholly unrelated to their current or future salary.

To assess the importance of these frictions – and how they may change when a scheme is closed – we estimate probit models of worker mobility. We start with the full sample of all workers in ASHE and then focus on the subset of workers in our sample of firms. In panel A of Table 3.6 we present estimates of the marginal effect of being in a DB scheme on annual mobility in those two samples. All regressions control for age and tenure effects and year and industry dummies. Looking at all workers, being a member of a DB scheme reduces the probability of exit from the firm by around 5 percentage points (compared to a mean exit rate of 24.5% for non-members). Panel B shows that the effects are significantly stronger within our sample of firms, with exit probabilities being about

9% lower for DB members. These results point to substantially lower mobility for those workers who are members of DB schemes. This is in line with earlier results obtained by McCormick and Hughes (1984) and Haverstick et al. (2010) in the US, and Disney and Emmerson (2002) in the UK. It is consistent with the lock-in hypothesis and helps us to understand how firms can impose some of the deficit costs on these workers.

What happens when the firm closes the scheme to future accruals? In the example above, the worker no longer gets the added benefit of the real wage growth effect because when a scheme is closed, final salaries are computed using the salary at closure date and subsequently uprated only by inflation. There is now no lock-in effect – the workers DB pension pay-out is identical whether the worker stays or leaves the firm - and we might expect to see increased mobility for DB members. Panel A of Table 3.7 shows that this is exactly what happens. The exit rate of those workers who were current members of the DB scheme rises by around 13% relative to those who were never members. This is similar in magnitude to the estimate of the lock-in from Table 3.6, therefore suggesting that post-closure, prior DB members regain a similar degree of mobility as all other workers who are not locked in to the firm by their pension benefits.

The increased mobility of DB members after scheme closure can be thought of as imposing a cost on firms, which mitigates the savings identified earlier. This is because greater mobility of workers may lead both to an outflow of firm-specific human capital and to additional hiring and training costs for the firm. The cost is likely to be higher the more experienced the departing workers are. Panel B of Table 3.7 assesses whether this is a likely concern for firms by comparing the mobility of high-wage and low-wage workers. We define high-wage workers as those whose average wage is above the median in the dataset. The results show that low-wage DB workers experience a relatively small degree of lock-in prior to scheme closure – 6% relative to non-DB low-wage workers – and their mobility

risers by the same 6% following scheme closure, relative to comparable non-DB workers. The effects are much larger and significant for high-wage DB workers, however. The exit rate of those workers rises by 18% after the scheme closes compared to similar non-DB workers. This suggests that the overall effect of removing pension-related frictions is driven to a greater extent by high earners and their pent-up demand for mobility. As the wages of those workers likely reflect high human capital – either general or firm-specific – their increased exit rates suggest a potentially significant cost for the firm.

3.6. Conclusions

This paper presents evidence on the impact of firm-specific cost shocks on workers. Using hand-collected data on the costs of funding historic pension liabilities, our analysis shows that firms have, to a limited extent, been able to transfer some of the cost of these liabilities onto workers in the form of lower wages. This burden sharing has been limited to those workers who have some exposure to the pension scheme, and even in this case, the firm bears the substantial share of the cost. However, firms have also increasingly chosen to close these pension schemes down in order to avoid incurring additional future liabilities. These closures have been associated with significant immediate savings for firms, as they contribute less to the replacement pension scheme. There is no evidence that workers are able to offset these losses by obtaining compensating wage increases.

These results imply frictions in the labour market that prevent the standard competitive results emerging. To explore these frictions in more detail, we examine the extent of job mobility. Importantly we show that workers who are currently members of a DB scheme have much lower exit rates than other workers. This effect is distinct from a tenure effect – though it is, of course, true that DB workers tend to have longer tenure with a firm. This helps us to understand why some of the burden of deficit payments can be

shifted onto these workers. However, when the firm closes down the DB scheme, we show that this “lock-in” vanishes – workers have the same exit rates regardless of former DB membership. This highlights a risk for the firm in closing down the scheme. Whilst immediate savings are made, there will be increased turnover. This is costly to the firm, and maybe more so given that the increased exits of DB workers who are now freed tend to come from the upper part of the wage distribution.

The paper adds to a growing literature on the importance of firm-specific shocks to wages and highlights yet again the shortcomings of the standard competitive model for the labour market. Frictions are important in understanding how workers are compensated and how firms deal with cost shocks.


3.7. Bibliography

- Abowd, J. M. (1989) “The Effect of Wage Bargains on the Stock Market Value of the Firm”, *American Economic Review*, 79, 774-800.
- Bunn, P., Mizen, P., and Smietanka, P. (2018) “Growing pension deficits and the expenditure decisions of UK companies”, Bank of England Working Paper No. 714.
- Chaudhry, N., Au Yong, H. H., and Veld. C. (2017) “How does the Funding Status of Defined Benefit Pension Plans Affect Investment Decisions of Firms in the United States”, *Journal of Business Finance and Accounting*, 44, 196-235.
- Disney, R., and Emmerson, C. (2002) “Choice of Pension Scheme and Job Mobility in the Britain”, IFS Working Paper No. 02/09.
- Garin, A., and Silverio, F. (2017) “How Does Firm Performance Affect Wages? Evidence from Idiosyncratic Export Shocks”, *mimeo*, Harvard University.
- Gruber, J. (1997) “The Incidence of Payroll Taxation: Evidence from Chile”, *Journal of Labor Economics*, 15, S72-101.
- Haverstick, K., Munnell, A. H., Sanzenbacher, G., and Soto, M. (2010) “Pension Type, Tenure, and Job Mobility”, *Journal of Pension Economics and Finance*, 9(4), 609-25.
- IAS 19 Employment Benefits. (2001) *International Financial Reporting Standards*. <http://eifrs.ifrs.org/eifrs/bnstandards/en/IAS19.pdf> (last accessed 22 March 2018).
- Kline, P., Petkova, N., Williams, H. and Zidar, O. (2017) “Who Profits from Patents? Rent-Sharing at Innovative Firms”, IRLE Working Paper No. 107-17, University of Berkeley.
- Liu, W., and Tonks, I. (2012) “Pension Funding Constraints and Corporate Expenditures”, *Oxford Bulletin of Economics and Statistics*, 75, 235-58.
- McCormick, B., and Hughes, G. (1984) “The Influence of Pensions on Job Mobility”, *Journal of Public Economics*, 23, 183-206.
- Munnell, A.H., and Soto, M. (2007). “Why Are Companies Freezing Their Pensions?”, Working Papers, Center for Retirement Research at Boston College No. 2007-22.
- OECD (2017) *Quarterly National Accounts*. <http://stats.ukdataservice.ac.uk/index.aspx?DatasetCode=QNA> (last accessed 1 August 2017).
- Phan, H. V., and Hegde, S. P. (2012) “Pension Contributions and Firm Performance: Evidence from Frozen Defined Benefit Plans”, *Financial Management*, 42, 373-411.
- Rauh, J. D. (2006) “Investment and Financing Constraints: Evidence from the Funding of Corporate Pension Plans”, *Journal of Finance*, 61, 33-71.

- Rauh, J. D., Stefanescu, I., and Zeldes, S. P. (2017) “Cost Saving and the Freezing of Corporate Pension Plans”, Stanford University Graduate School of Business Research Paper No. 16-4.
- Van Reenen, J. (1996) “The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies”, *Quarterly Journal of Economics*, 111(1), 195-226.


3.8. Figures and tables

FIGURE 3.1: ACTIVE MEMBERS OF DEFINED BENEFIT PENSION SCHEMES, 1979-2015




Notes: The chart plots the total number of workers who are active members of DB pension schemes in the UK, by sector and scheme status. Source: Pension Regulator.

FIGURE 3.2: AGGREGATE NORMAL AND SPECIAL PENSION CONTRIBUTIONS, 1992-2015




Notes: The chart plots the aggregate estimates of employers' normal (regular) and special (deficit) contributions to pension schemes from the UK national accounts. Source: Office for National Statistics.

FIGURE 3.3: COMPONENTS OF TOTAL COMPENSATION OF EMPLOYEES IN THE UK, 1972-2015




Notes: The chart plots the breakdown of aggregate compensation of employees from the UK national accounts. Source: Office for National Statistics.

FIGURE 3.4: ECONOMY-WIDE VS SAMPLE DB CONTRIBUTIONS, 2003-2015




Notes: The chart plots the aggregate estimates of employers' DB contributions to pension schemes from the UK national accounts (UK Aggregate) and the total for our sample of UK-listed firms (Our Sample). Source: Office for National Statistics.

FIGURE 3.5: DEFICIT CONTRIBUTIONS: OUR DEFINITION VS. COMPANY REPORTS, 2002-2015




Notes: The chart compares our calculation of deficit DB contributions (total employer DB contributions – current service cost) to the deficit contributions as reported by the firms themselves, for the 111 firms in our dataset for which we are able to extract both measures from the annual reports in at least one of the years.

FIGURE 3.6: DISTRIBUTION OF DEFICIT PAYMENTS PER WORKER IN THE SAMPLE




Notes: The chart shows the distribution of the pension deficit measure, `db_emp`, used in the regressions. The few values below zero correspond to cases of firms underfunding their DB schemes. 66% of those cases occur before the financial crisis (2007 or earlier) in the data, when relatively more schemes were in surplus than in the later period, and some firms were allowed to reduce or stop their contributions.

FIGURE 3.7: SCHEME CLOSURE AND EMPLOYER CONTRIBUTIONS



Notes: This chart plots the evolution of employer contributions for the period between four years prior to DB scheme closure and four years subsequent to DB scheme closure, separately for exposed DB members and all other workers. Exposed DB members are defined as those whose employer reported a positive contribution in the year the scheme closed. Blue and red dots relating to the same period have been shifted horizontally for clarity. The default category are employees at firms that did not close their DB schemes.

FIGURE 3.8: SCHEME CLOSURE AND WAGES



Notes: This chart plots the evolution of wages for the period between four years prior to DB scheme closure and four years subsequent to DB scheme closure, separately for exposed DB members and all other workers. Exposed DB members are defined as those whose employer reported a positive contribution in the year the scheme closed. Blue and red dots relating to the same period have been shifted horizontally for clarity. The default category are employees at firms that did not close their DB schemes.

TABLE 3.1: SUMMARY STATISTICS

	Mean	Median	Standard Deviation
A) Firm Sample			
Employment	10,361	2,096	27,111
Revenue (£m)	2,209	327	6,227
Wage Bill (£m)	297	55	909
DB Plan (0/1)	0.72	1.00	0.45
DB Pension Liabilities (£m)	2,094	391	5,380
DB Net Accounting Deficit (£m)	-159	-25	698
DB Deficit Payments per Employee (£ 000)	2.09	0.84	3.73
DB Member %	29.86	22.22	30.31
Sample Size (firm-year obs.)		2,638	
Number of Firms		330	
B) Worker Sample			
Hourly Wage (£)	13.70	10.19	11.83
Weekly Wage (£)	479	376	441.26
DB Employer Contribution Rate (%)	11.79	10.41	10.48
Tenure (years)	8.80	6.00	8.84
DB Member (0/1)	0.27	0.00	0.45
Sample Size (worker-year obs.)		197,748	
Number of Workers		55,202	

Notes: Panel A (Firm Sample): data were hand-collected from the annual reports of the 450 listed firms in the sampling frame described in the main text, publicly available from Companies House. The data cover the period 2001-2015. Panel B (Worker Sample): data are from the Annual Survey of Hours and Earnings, 2002-2016. Values are shown in nominal terms.

TABLE 3.2: PENSION DEFICITS AND WORKERS' WAGES

	Hourly Wage			Weekly Wage	
	(1)	(2)	(3)	(4)	(5)
db_emp(-1)	-0.085* (0.050)	-0.128** (0.056)	-0.127** (0.056)	-0.087 (0.060)	-0.133** (0.069)
db_emp(-2)	-0.151* (0.082)	-0.178* (0.092)	-0.171* (0.092)	-0.148 (0.091)	-0.175* (0.096)
Σ db_emp	-0.237* (0.123)	-0.306** (0.142)	-0.298** (0.144)	-0.235* (0.134)	-0.308** (0.158)
Individual FE	x				
Year FE	x	x	x	x	x
5-digit Industry FE	x				
1-digit Industry * Year FE		x	x	x	x
Firm Performance Controls				x	
Firm FE		x			
Match FE			x	x	x
Sample Size	197,748	197,748	197,748	158,396	197,748

Notes: The dependent variable is the log of hourly wages, unless specified otherwise in the column heading. Reported coefficient estimates are multiplied by 100. All regressions include age, age squared and tenure. Standard errors are clustered at the firm level. ***, **, and * indicate significance at the 1%, 5% and 10% level respectively.

TABLE 3.3: PENSION DEFICITS AND DB MEMBERSHIP

	(1)	(2)	(3)	(4)	(5)	(6)
$\sum db_emp$ (DB Member)	-0.410** (0.163)	-0.458** (0.179)	-0.380*** (0.145)			
$\sum db_emp$ (Non-DB Member)	-0.048 (0.093)	-0.025 (0.091)	0.143 (0.205)			
$\sum db_emp$ (Large Active Membership)				-0.355** (0.151)	-0.402** (0.163)	-0.240** (0.116)
$\sum db_emp$ (Small Active Membership)				-0.114 (0.143)	-0.106 (0.143)	-0.180 (0.328)
Estimated Sharing Rate	9.8%	11.0%	9.1%	8.5%	9.6%	5.8%
Individual FE	x			x		
Year FE	x	x	x	x	x	x
5-digit Industry FE	x			x		
1-digit Industry x Year FE	x	x	x	x	x	x
Firm Performance Controls			x			x
Firm FE						
Match FE		x	x		x	x
Sample Size	197,748	197,748	158,413	197,748	197,748	158,413

Notes: The dependent variable is the log of hourly earnings. All regressions include age, age squared and tenure. Each row reports the sum of the coefficients (multiplied by 100) for different DB membership at the individual or firm level on the pension deficit measure, which has two lags included in all specifications. Standard errors are clustered at the firm-level. The estimated sharing rate is calculated by multiplying the estimate of $\sum db_emp$ for DB members or for workers at firms with large active membership by the mean hourly wage of £13.70, mean weekly hours of 33.7 and 52 weeks in a year (calculated from the sample of 197,748), divided by 1000 since annual deficit payments per worker are expressed in £ thousands. ***, **, and * indicate significance at the 1%, 5% and 10% level respectively.

TABLE 3.4: SCHEME CLOSURE AND EMPLOYER PENSION CONTRIBUTIONS

	All	Employer rate not zero	DB rate positive in year of close	Scheme closed in valuation year
	(1)	(2)	(3)	(4)
Scheme Closed (Exposed DB Members)	-7.679** (3.556)	-7.407** (3.515)	-6.542*** (2.103)	-14.036*** (2.093)
Scheme Closed (All Other Workers)	0.691 (1.051)	1.000 (0.919)	-2.735 (2.385)	-2.820* (1.542)
Year FE	x	x	x	x
1-Digit Industry * Year FE	x	x	x	x
Match FE	x	x	x	x
Sample Size	112,497	106,932	112,497	95,090

Notes: The dependent variable is the employer pension contribution rate (in percentage points). The table reports the comparison between employer contributions post- and pre-closure for exposed DB members and all other workers, as defined in the main text. All regressions include age, age squared and tenure. Standard errors are clustered at the firm-level. ***, **, and * indicate significance at the 1%, 5% and 10% level respectively.

TABLE 3.5: SCHEME CLOSURE AND WAGES

	All	Employer rate not zero	DB rate positive in year of close	Scheme closed in valuation year
	(1)	(2)	(3)	(4)
Scheme Closed (Exposed DB Members)	-0.017 (0.017)	-0.013 (0.015)	-0.017 (0.021)	0.008 (0.036)
Scheme Closed (All Other Workers)	-0.001 (0.012)	0.001 (0.009)	-0.004 (0.012)	-0.007 (0.012)
Year FE	x	x	x	x
1-Digit Industry * Year FE	x	x	x	x
Match FE	x	x	x	x
Sample Size	197,748	106,932	197,748	160,509

Notes: The dependent variable is the log of hourly wages. The table reports the comparison between employer contributions post- and pre-closure for exposed DB members and all other workers, as defined in the main text. All regressions include age, age squared and tenure. Standard errors are clustered at the firm-level. ***, **, and * indicate significance at the 1%, 5% and 10% level respectively.

TABLE 3.6: PENSION SCHEME MEMBERSHIP AND WORKER MOBILITY

A) All Workers in ASHE	
DB Member	0.196
Not a DB Member	0.243
Sample Size	1,035,302
B) Workers in Sample Firms	
DB Member	0.180
Not a DB Member	0.280
Sample Size	233,664

Notes: This table reports the probabilities of worker exit based on DB pension membership status. The results are based on a probit regression of a 0/1 indicator of whether the worker leaves the firm in the next period on a scheme membership indicator, as well as controls for age, age squared, gender, tenure, tenure squared, year, 2-digit industry, and firm size measured as the sales quartile.

TABLE 3.7: PENSION SCHEME CLOSURE AND WORKER MOBILITY

	(1) Pr(Exit Scheme Open)	(2) Pr(Exit Scheme Closes)	(2) - (1) Effect of Scheme Closure on Pr(Exit)
A) All Workers in Sample Firms			
(a) DB Member	0.189*** (0.002)	0.391*** (0.022)	0.201*** (0.022)
(b) Not a DB Member	0.275*** (0.002)	0.346*** (0.007)	0.070*** (0.007)
Difference-in-Difference (a-b)			0.131*** (0.023)
B) Workers in Sample Firms by Position in the Wage Distribution			
<i>(i) Workers with Wage ≤ Median</i>			
(a) DB Member	0.209*** (0.004)	0.341*** (0.032)	0.132*** (0.032)
(b) Not a DB Member	0.276*** (0.002)	0.347*** (0.008)	0.071*** (0.008)
Difference-in-Difference (a-b)			0.061* (0.033)
<i>(ii) Workers with Wage > Median</i>			
(a) DB Member	0.180*** (0.003)	0.425*** (0.029)	0.245*** (0.029)
(b) Not a DB Member	0.283*** (0.003)	0.352*** (0.015)	0.069*** (0.015)
Difference-in-Difference (a-b)			0.177*** (0.033)

Notes: This table reports the marginal effect of DB scheme closure on worker exit based on DB pension membership status. Column (1) shows the results from a probit regression of a 0/1 indicator of whether the worker leaves the firm in the next period on scheme membership and scheme closure indicators and their interaction, as well as controls for age, age squared, gender, tenure, tenure squared, year, and 2-digit industry, and firm size measured as the sales quartile. Low-wage workers are defined as those who are, on average, below or at the median of the annual wage distribution in ASHE, while high-wage workers are those who are on average above the median. Predicted exit probabilities and marginal effects are calculated at the means of all regressors. Sample size is 173,444. Deferred members and post-closure periods are excluded from the sample. ***, **, and * indicate significance at the 1%, 5% and 10% level, respectively.

3.9. Appendices

Appendix A – Data

This appendix shows a breakdown of our sample by year and workers' pension membership status to illustrate that it reflects the decline in aggregate defined-benefit pension scheme membership in the UK over the same period. The appendix also describes the main variables used in the empirical analysis.

Table 3.A1 shows the number of workers and firms in the sample by year from 2002 to 2016. Since the empirical model uses two lags of the pension deficit variable, our analysis uses worker observations starting in 2004. We observe a decline in DB members over time, both in absolute and percentage terms. This is in line with the aggregate trends shown earlier in Figure 3.1. There is a corresponding increase in the number and proportion of DC members, particularly from 2013 onwards, when new legislation on pension auto-enrolment led to the employees of some UK firms being automatically subscribed to occupational pension schemes by their employers. Overall, therefore, our sample reflects important aspects of the evolution of the aggregate UK pension landscape.

Tables 3.A2 and 3.A3 describe the main raw and derived variables used in our empirical analysis. Firm-level data were collected from the annual reports of firms or sourced from the Annual Respondents Database (ARD, 2002-2007) and the Annual Business Survey (ABS, 2008-2016). Data on the compensation, pension membership, and personal characteristics of workers were sourced from the Annual Survey of Hours and Earnings (ASHE). We matched firm and worker data by mapping the Dun and Bradstreet codes of the firms in our sample to the enterprise group reference number (*wowent*). We then mapped the *wowent* variable to enterprise reference numbers (*entref*). This allowed us to link our pension data to firms ('enterprises') in ARD and ABS, as well as to the individuals working for those firms and captured in ASHE.

TABLE 3.A1: SAMPLE OBSERVATIONS BY YEAR

Year	Number of Firms in the Sample	Number of Workers in the Sample			
		DB Member	Not a DB Member		Total
			DC Member	No Pension	
2002	50				
2003	61				
2004	191	4,878	2,975	3,036	10,889
2005	233	7,623	3,176	6,828	17,627
2006	237	5,931	4,533	5,980	16,444
2007	209	4,994	3,998	5,832	14,824
2008	205	4,371	4,480	6,499	15,350
2009	200	4,535	4,235	7,175	15,945
2010	194	4,345	4,543	7,217	16,105
2011	187	4,013	4,640	7,629	16,282
2012	185	3,007	4,929	7,332	15,268
2013	176	2,962	6,666	5,194	14,822
2014	174	2,751	8,890	3,674	15,315
2015	173	2,714	8,654	3,256	14,624
2016	163	2,206	8,167	2,833	13,206
Total		54,330	69,886	72,485	196,701

Notes: This table reports the number of worker observations in the sample broken down by pension membership status and year, as well as the annual number of firms that employ these workers. There are 1,047 additional worker-year observations with no data on pension membership status available in ASHE, giving a total sample size of 197,748 observations. There are 330 unique firms and 55,202 unique workers in the sample.

TABLE 3.A2: DESCRIPTION OF KEY FIRM-LEVEL VARIABLES

Variable	Description	Source
Total employer contributions	Employer contributions to the firm's UK DB pension schemes	Annual reports
Current service cost	Increase in the present value of a defined benefit obligation resulting from employee service in the current period	Annual reports
DB assets	Accounting value of assets in the firm's UK DB pension schemes	Annual reports
DB liabilities	Accounting value of assets in the firm's UK DB pension schemes	Annual reports
DB deficit/(surplus)	DB liabilities minus DB assets	Annual reports
Scheme closure date	Year when the firm closed its last DB pension scheme	Annual reports
Deficit payments	Total employer contributions – Current service cost	Annual reports
Sales	Annual sales (<i>turnover</i> variable in ARD, <i>wq550</i> variable in ABS) matched to employees based on the enterprise reference number (<i>entref</i>)	ARD/ABS
Profitability	Gross value added (sales minus the cost of intermediate inputs other than employment costs; <i>gva_fc</i> variable in ARD, <i>wq612</i> variable in ABS) divided by sales	ARD/ABS
Industry	2-digit SIC 2003 code corresponding to the enterprise whether the employee works (<i>sic03</i> variable). For 2012-2015, the SIC 2003 code is derived from the SIC 2007 code (<i>sic07</i> variable) following the mapping based on gross output developed by Richard Harris and available to researchers working on UK survey data through the UK Data Service	ASHE

TABLE 3.A3: DESCRIPTION OF KEY WORKER-LEVEL VARIABLES

Variable	Description	Source
Hourly wage	Gross hourly wage including basic pay, incentive pay, shift and premium payments, and overtime	ASHE
Weekly wage	Gross weekly earnings including basic pay, incentive pay, shift and premium payments, and overtime	ASHE
Employer contributions	Employer contributions to the worker's occupational pension scheme(s)	ASHE
Age	Employee age	ASHE
Tenure	Difference between the year when the worker is observed (<i>year</i> variable) and the year when the employee joined at the firm (derived from the <i>empsta</i> variable, which gives the month and year of the start date)	ASHE
Pension membership	Type of pension benefit, classified into DB, DC or no pension based on the <i>pens</i> variable for 2002-2003 and the <i>tpen</i> variable from 2004 onwards	ASHE

Appendix B – Additional tables

TABLE 3.B1: TESTING THE EFFECT OF THE 3RD LAG OF DEFICIT PAYMENTS ON WAGES

	(1)
db_emp(-1)	-0.078 (0.072)
db_emp(-2)	-0.150* (0.091)
db_emp(-3)	-0.028 (0.118)
$\sum db_emp$	-0.255 (0.177)
Individual FE	
Year FE	x
5-digit Industry FE	
1-digit Industry * Year FE	x
Firm Performance Controls	x
Firm FE	
Match FE	x
Sample Size	145,445

Notes: This table shows the results of testing the significance of the third lag of the pension deficit measure, db_emp. The dependent variable is the log of hourly earnings. Reported coefficient estimates are multiplied by 100. All regressions include age, age squared and tenure. Standard errors are clustered at the firm level.

* indicates significance at the 10% level.

TABLE 3.B2: PENSION DEFICITS AND DB MEMBERSHIP USING LAGGED EMPLOYMENT

	(1)	(2)	(3)
Σdb_emp_lag	0.079 (0.086)		
Σdb_emp_lag (DB Member)		-0.381*** (0.145)	
Σdb_emp_lag (Non-DB Member)		0.140 (0.206)	
Σdb_emp_lag (Large Active Membership)			-0.242** (0.116)
Σdb_emp_lag (Small Active Membership)			-0.179 (0.382)
Estimated Sharing Rate		9.1%	5.8%
Year FE	x	x	x
1-digit Industry x Year FE	x	x	x
Firm Performance Controls	x	x	x
Match FE	x	x	x
Sample Size	148,927	158,396	158,396

Notes: The dependent variable is the log of hourly earnings. All regressions include age, age squared and tenure. Each row reports the sum of the coefficients (multiplied by 100) for different DB membership on the pension deficit measure, db_emp_lag , which has two lags included in all specifications. Here, the pension deficit measure is calculated as deficit payments divided by prior period's employment (rather than starting employment like in the main text). All standard errors are clustered at the firm-level. The estimated sharing rate is based on a mean hourly wage of £13.70 (calculated from the sample of 197,748). ***, **, and * indicate significance at the 1%, 5% and 10% level, respectively.

TABLE 3.B3: SCHEME CLOSURE AND EMPLOYEE COMPENSATION

Year relative to scheme closure year (t)	Marginal effect of time period on employer contributions		Marginal effect of time period on wages	
	Exposed DB members	Not exposed	Exposed DB members	Not exposed
t-4	2.418 (1.730)	2.579 (1.207)	-0.005 (0.023)	-0.016 (0.011)
t-3	2.723 (1.170)	1.474 (1.326)	0.067 (0.018)	0.028 (0.024)
t-2	2.718 (1.653)	2.909 (1.302)	-0.002 (0.029)	-0.020 (0.013)
t-1	2.021 (1.509)	2.782 (1.375)	-0.014 (0.020)	-0.024 (0.009)
t	1.589 (1.308)	5.609 (4.297)	-0.005 (0.023)	-0.009 (0.012)
t+1	-3.181 (2.015)	1.275 (1.590)	-0.019 (0.031)	-0.017 (0.017)
t+2	-3.990 (1.321)	0.128 (0.835)	0.002 (0.038)	-0.010 (0.015)
t+3	-4.038 (2.350)	1.306 (1.511)	-0.010 (0.029)	-0.018 (0.014)
t+4	-4.306 (2.654)	0.441 (1.307)	0.061 (0.070)	0.009 (0.016)

Notes: This table shows the evolution of employer contributions and wages for the period between four years prior to DB scheme closure and four years subsequent to DB scheme closure. It reports the marginal effects of the time periods relative to the year of DB scheme closure (t) on employer contributions and wages for exposed DB members and all other workers. Exposed DB members are defined as those whose employer reported a positive contribution in the year the scheme closed. Regressions include year, 1-digit industry*year, and match fixed effects. The default category are employees at firms that did not close their DB schemes. The sample size is 112,497 for employer contributions and 197,555 for wages.

Conclusion

This thesis investigates the impact of firm characteristics and firm-level events on wages and employment, using data on firms and individuals from the UK. The first chapter finds that on average, individuals who join young firms receive a small wage premium at the outset but can expect lower subsequent wage growth than individuals who join mature firms. Well-paid and stable jobs at young firms do exist, but they are relatively rare. These results contribute to a nuanced assessment of the job-creating capabilities of young firms, which have attracted the interest of policymakers in many countries.

The second chapter finds that firms with a higher market share and greater capital intensity of production share a smaller proportion of their value added with workers than less “mighty” firms do. The results suggest that firm-level drivers may be an important explanation behind the decline in the aggregate labour share observed in many economies around the world since the 1970s.

Finally, the third chapter finds that firms are able to share the burden of unexpected costs related to legacy defined-benefit pension plans with plan members, due to the high attachment of workers to firms that such pension benefits generate. Firms that close their defined-benefit pension plans are able to make savings on payroll costs. However, such closures reduce the frictions that workers face and thus carry an implicit cost for the firm. This highlights the importance of labour market frictions for understanding how employees are compensated and how firms respond to shocks.

These findings generate several interesting questions for future research. One such question is how the changing nature of the universe of start-ups has contributed to the slowdown in productivity growth in the UK over the last decade. Productivity growth has important implications for the evolution of wages and living standards. While many

explanations have been proposed for the recent slowdown, the role of firm entry and the changing characteristics of young firms remain unexplored. Investigating the underlying reasons for the decline in the productivity and survival prospects of young firms may generate tangible implications for policymakers interested in understanding, and reviving, aggregate productivity growth.

The analysis of firm-level labour shares provides a basis for new analyses that relate my results back to macroeconomic trends. One direction for future research may be to decompose the historical movements in the aggregate labour share into changes at the firm level and changes in the distribution of firms across capital intensity and market share categories within sectors. The analysis could also be replicated using firm-level data from countries other than the UK and expanded to cover a longer time period that would coincide with more dramatic labour share declines. Moreover, causal links between capital intensity, market power, and the labour share can also be explored further, perhaps by identifying natural experiments that provide quasi-exogenous variation in access to capital across firms or in market structure across sectors.

Finally, the analysis of firms' responses to unexpected pension deficit payments can be extended to other sources of idiosyncratic cost changes, either positive or negative. Such work can enhance our understanding of how firms respond to shocks, and of the extent to which labour market frictions hinder the transmission of such shocks to wages.

