



BANK OF ENGLAND

Working Paper No. 495

The productivity puzzle: a firm-level investigation into employment behaviour and resource allocation over the crisis

Alina Barnett, Adrian Chiu, Jeremy Franklin and
María Sebastiá-Barriel

April 2014

Working papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee.



BANK OF ENGLAND

Working Paper No. 495

The productivity puzzle: a firm-level investigation into employment behaviour and resource allocation over the crisis

Alina Barnett,⁽¹⁾ Adrian Chiu,⁽²⁾ Jeremy Franklin⁽³⁾ and María Sebastiá-Barriel⁽⁴⁾

Abstract

Labour productivity in the United Kingdom has been exceptionally weak since the 2007/08 financial crisis. This paper uses firm-level data from the Office for National Statistics Annual Business Survey and the Inter-Departmental Business Register to better understand the nature of this weakness. Overall, our findings are consistent with existing literature which finds that within-firm productivity growth tends to be procyclical and emphasises the importance of the reallocation of resources between firms and sectors for productivity growth. More specifically, we find that up until 2011 there was a doubling in the proportion of firms with shrinking output and flat employment. This suggests that firms were able to respond flexibly to weak demand conditions by retaining staff at the expense of measured productivity, suggestive of an opening up of spare capacity within firms. However, the strength of recent hiring behaviour since 2012 means that this is now likely to be less of a factor. The lack of labour shedding, together with a low firm exit rate, is also indicative of low levels of resource reallocation between firms and sectors. To assess the importance of this to aggregate productivity growth we apply the method used by Baily, Bartelsman and Haltiwanger. We find that reallocation between firms (in terms of both the movement of labour and firm entry and exit) contributed significantly to aggregate productivity growth before the crisis, but its contribution fell substantially after. In fact, around one third of the productivity slowdown after 2007 can be attributed to slower reallocation of resources. The extent to which reduced factor reallocation, and so the weakness in productivity growth, persists remains a key question for the economic outlook.

Key words: Productivity growth, long-run growth, resource reallocation, entry, exit, financial crisis.

JEL classification: L11, O47, E32.

(1) Bank of England. Email: alina.barnett@bankofengland.co.uk

(3) Bank of England. Email: jeremy.franklin@bankofengland.co.uk

(2) Bank of England. Email: adrian.chiu@bankofengland.co.uk

(4) Bank of England. Email: maria.barriel@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We are indebted to the Virtual Microdata Laboratory team at the Office for National Statistics (ONS) for providing and supporting access to their comprehensive firm-level database, especially Steve Bond, Felix Ritchie and James Stainer. We are also grateful to Sandra Batten, Jamie Bell, Stuart Berry, Ben Broadbent, Alan Castle, Mark Franklin, Bob Gilhooly, Richard Harris, Mark Hart, Jonathan Haskel, Jenni Greenslade, Neal Hatch, Helen Miller, Katharine Neiss, John Van Reenen, Jane Walsh, and Martin Weale for useful comments and discussions. This work contains statistical data from the ONS, which is Crown Copyright. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research data sets that may not exactly reproduce National Statistics aggregates. This paper was finalised on 26 March 2014.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH

Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

Summary

Labour productivity growth in the United Kingdom has been exceptionally weak since the 2007/08 financial crisis and currently lies around 14 percentage points below the level implied by its pre-crisis trend growth rate. This phenomenon is commonly referred to as the UK ‘productivity puzzle’. Such a prolonged period of weakness in labour productivity stands out from historical and international experiences.

This paper uses firm-level data collected by the Office for National Statistics (ONS) to improve our understanding about the drivers of the weakness in UK labour productivity. This analysis only covers the period to 2011, so cannot shed light on the strength in UK employment in 2012 and 2013.

There are two main objectives. First, to set out some stylised facts about productivity across both time and firm dimensions. Within this part we also discuss how representative our results are for the UK economy as a whole by comparing employment and GVA in our sample of firms to ONS aggregate statistics. We observe that the productivity experience across firms of different sizes has been varied. Labour productivity fell more sharply for small and medium-sized enterprises (SMEs) at the onset of the crisis than for large firms, but by 2011 the weakness in productivity relative to its pre-crisis trend appears relatively evenly distributed across SMEs and large firms. Second, we examine the changing role of resource reallocation on UK productivity growth before and after the financial crisis.

There are two key findings that stand out. First, we find that a large part of the persistent weakness in productivity can be accounted for by the fact that the proportion of firms with shrinking output and flat employment doubled from 11% in 2005-07 to 22% in 2011. At any point in the economic cycle we see some firms who are growing and hiring more workers, while other firms are shrinking and reducing employment. At the onset of the financial crisis there was a significant decline in the proportion of firms that were growing and hiring, and a rise in the number of firms that were shrinking and firing. But by 2011, a large concentration of firms emerged that had shrinking output but no change in employment. This suggests that firms were able to respond flexibly to weak demand conditions by retaining staff at the expense of measured productivity, suggestive of an opening up of spare capacity within firms. This result implies a more *temporary* or *cyclical* explanation of the productivity slowdown in the United Kingdom, as these firms may have been well placed to increase production quickly without hiring when demand conditions started to strengthen. However, the strength of recent hiring behaviour since 2012 implies that this may now be less of a factor.

The second key finding relates to the role resource allocation played in the slowdown of labour productivity growth. Reallocation is the process through which factors of production move from lower to higher productivity firms, helping to take advantage of market expansion opportunities and generate aggregate increases in productivity. We find that labour reallocation, which includes movements in labour brought about by company formation and dissolution, and within-firm productivity improvements were equally important in driving productivity growth between 2002 and 2007. However since the crisis, the role of reallocation fell significantly while the

contribution of within-firm productivity to aggregate productivity growth turned negative. Taken together, we find that a third of the shortfall in labour productivity since the crisis can be attributed to slower resource allocation across firms. The rest is due to productivity weakness within firms. This result is indicative of a more *persistent* interpretation of the UK productivity slowdown. There may be a number of factors that have slowed the reallocation process; for example a disrupted financial sector or heightened uncertainty about the economic environment. Since our data are at the firm and not plant level, these within-firm changes will not include any reallocation effects across individual plants that sit within particular firm entities and may, therefore, underestimate the scale of these reallocation effects.

We believe that the effect of reallocation on measured productivity may have been low in part because a greater proportion of firms facing difficulties have managed to survive the latest recession compared to previous ones. As discussed in the Bank of England's August 2013 *Inflation Report*, the number of company liquidations has remained unusually low this recession given developments in output. There may be several factors that may have helped firms survive, for example the low level of Bank Rate, weak real wage growth and any forms of loan forbearance (the extent of SME forbearance is discussed in the 2013 Q4 *Quarterly Bulletin*). To try to understand the likely magnitude of this effect, we develop a scenario to examine what might have happened to measured labour productivity should firm deaths have increased to a level more consistent with the 1990s recession. Because our data set only goes back to 1997, and therefore does not include previous recessions, we use ONS aggregate statistics to inform our counterfactual exercise. We find that the unusually low level of business failure is likely to have materially lowered *measured* labour productivity. Nonetheless, lower business failure, and the resultant lower unemployment, probably meant that the loss to GDP and general welfare associated with the financial crisis was smaller than it otherwise would have been.

Overall, our key findings suggest that the slowdown in UK productivity is likely to have coincided with an opening of spare capacity at the onset of the recession, as firms reacted flexibly to the weakness in demand by retaining staff. However, the strength in aggregate employment since 2012 implies that this may now be less of a factor. We also find that reallocation between firms (in terms of both the movement of labour and firm entry and exit) contributed significantly to aggregate productivity growth before the crisis, but its contribution fell substantially after. The speed at which labour productivity is able to grow in the short to medium term may be limited by the extent to which impaired allocation of resources across the economy continues to be a binding constraint. But exactly how companies and resources will respond as demand recovers remain key questions for the economic outlook.

1 Introduction

Labour productivity growth in the UK has been exceptionally weak since the 2007/8 financial crisis. Whole-economy output per worker fell by 5% between 2007 and 2009, and has grown by an average of 0.3% per year since. This compares to an average annual growth rate of 2.6% in the decade up to 2007. This recent weakness in labour productivity stands out from historical experiences and is unprecedented during peace times (Chart 1). This also stands out from the experiences in several other countries, particularly the US (see for example Hughes and Saleheen, 2012 and ONS, 2012). Whole-economy output per worker in 2013Q3 lies around 14pp below the level implied by its pre-crisis trend growth rate. This is what is commonly termed the UK ‘productivity puzzle’.

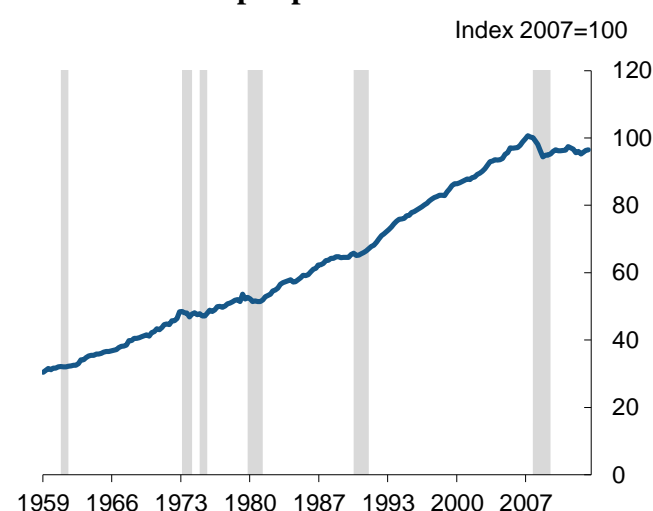
The evolution of labour productivity is an important consideration for monetary policy, as it provides an important indicator of an economy’s ability to supply goods and services without putting upward pressure on wage and price inflation. Indeed, a variety of explanations for the productivity puzzle have emerged in recent years which have been at the forefront of discussions by the Monetary Policy Committee (for example, see the January and February 2013 MPC minutes).

One candidate explanation is that efficient resource allocation across sectors and firms of the economy has been impaired since the crisis. A growing body of evidence, particularly for the US, highlights the importance of efficient labour reallocation and firm restructuring as a source of long-run productivity growth (amongst others see Bartelsman, Haltiwanger and Scarpetta (2013)).

This paper uses firm-level data collected by the Office for National Statistics (ONS) to better understand the drivers of UK labour productivity. There are two main objectives for this research. First, we document the heterogeneity of labour productivity dynamics across time and firms. Within this part we also discuss how representative our results are for the whole UK economy by comparing sample employment and GVA to the ONS aggregate statistics. We find that broad trends are similar, thus we are confident that we are able to make inferences across firms based on the sample dataset that we have compiled. Our analysis here suggests the fall in productivity appears broad based across firm sizes, and at least partly related to higher labour retention behaviour by firms over the period to 2011.

Second, we examine the role of resource reallocation on UK productivity growth before and after the 2007 financial crisis. We find that although reallocation of labour across firms,

Chart 1: UK output per worker 1959-2013



Notes: Grey bars indicate technical recessions (two periods of consecutive negative quarterly GDP growth). Source: ONS.

including new company births and deaths, was an important driver of aggregate productivity growth before the crisis, this has halved in the period after the crisis.

This paper builds on existing work in this area and provides a comprehensive analysis of firm-level productivity across the UK economy. Previous studies on this topic have tended to either cover a narrow part of the economy, like manufacturing only, or analysed earlier time periods (see for example Harris and Moffat, 2013 and Disney, Haskel and Heden, 2003), or used alternative methods to answer a more specific question (Riley, Rosazza Bondibene and Young, 2013, Crawford, Jin and Simpson, 2013), or provided detailed statistics (ONS, 2012).

The paper is structured into five parts. First, we discuss how our work fits with the growing body of work that examines the UK productivity puzzle in general, and that which analyses the role of resource reallocation on productivity growth in particular. Second, we provide a description of the data and show how this compares with the ONS aggregate statistics for the whole UK economy. The subsequent section discusses the degree of heterogeneity across firms over time. Next we show how hiring and firing behaviour across firms changed during the recent recession. We then turn to decomposing UK aggregate productivity growth into the contribution from reallocation – specifically, the contribution from firm births, deaths, market share changes – and within-firm changes. We also conduct a scenario analysis to determine the extent to which relatively low firm deaths in this recession may have contributed to the weakness we have seen in measured labour productivity. Last, we conclude with the key findings and the policy implications.

2 Relevance to the ‘Productivity Puzzle’ and existing literature

In this section we first review the existing work that has been done on the UK ‘Productivity Puzzle’. We then discuss the role of resource reallocation on productivity growth and explain how our paper contributes to this growing area of research.

2.1 The UK productivity puzzle

There are two main competing hypotheses for the existence of the puzzle which have very different and important implications for monetary policy.

The first hypothesis argues that the weakness in productivity is *temporary or cyclical*, reflecting transient changes in demand conditions which may simply reverse as and when the economy recovers. In this scenario, if firms face costs to adjusting their labour force they may choose to hold on to labour in order to retain their skills and experience for when the economy recovers – what is known as ‘labour hoarding’¹ (Oi, 1962). In this case, the weakness in productivity we have seen is likely to be temporary and does not reflect any technological or supply constraints but rather an increase in spare capacity within firms. Weak real wages since the crisis are likely

¹ Oi (1962) is the classical reference on labour-hoarding. He motivates the existence of labour hoarding through the adjustment costs firms face when adjusting skilled labour force. Others like Fay and Medoff (1985) mention other reasons such as the value of retaining a skilled labour stock in anticipation of recovery, as well as contractual commitments and the adverse implications of labor-force adjustment for employee morale.

to have encouraged firms to switch from more capital intensive to more labour intensive forms of production. Blundell, Crawford and Jin (2013) indeed find that flexible wages and increased labour supply are likely to have affected aggregate productivity. Goodridge, Haskel and Wallis (2013) point out that labour hoarding amongst skilled workers may also be a reflection of the economy's inability to measure intangible output properly. In addition, if businesses divert additional resources to activities dedicated to winning contracts or find it harder to win work, this may also lead temporarily to weak measured productivity.

The second hypothesis argues that the weakness in productivity we have seen is more *persistent* or *structural*, reflecting shocks that have a persistent effect on the productive capacity of the UK. Negative shocks to the availability of credit to UK companies could be such an example. Indeed Oulton and Sebastia-Barriel (2013) find that financial crises tend to reduce the long-run level of productivity. As noted in the Bank of England's November 2013 *Inflation Report*, there are two further pieces of evidence in support of this hypothesis. First, there is very little spare capacity within firms reported by the business surveys. The second is the strength in employment. The level of employment was 3% higher in 2013Q1 than 2010Q1 meaning total employment has increased by around 900,000 people during this period. If firms are increasing employment, this suggests that firms are unable to meet changes in demand with existing staff levels – suggestive of little spare capacity in firms. Now it may well be that any factors inhibiting productivity growth subside in future allowing productivity to recover. However, at least in the short to medium term a more structural interpretation of the productivity puzzle implies less scope for monetary policy to stimulate demand growth without excessive inflationary pressure.

Now it could well be that the behaviour of productivity in the UK cannot be easily ascribed to one of the above categories. In fact, the narrative behind the fall in labour productivity could fall anywhere in the wide spectrum between these two hypotheses.

In this paper, we contribute to this body of work by bringing new evidence on the productivity experience of different types of firms and sectors before and after the recent financial crisis. We find that a large part of the recent weakness can be accounted for by a doubling in the proportion of firms with shrinking output and flat employment. Indeed, we find that the employment behaviour of firms may have changed since the crisis which can explain part of the weakness in productivity.

2.2 *Reallocation, restructuring and productivity growth*

Another economic literature that our study contributes to is on the role of resource allocation in driving aggregate productivity growth. The intuition here is that reallocation is the process through which factors of production move from low to high productivity firms, helping to take advantage of market expansion opportunities and generate aggregate increases in productivity. These studies typically decompose aggregate productivity growth into a within-firm growth component and a reallocation of factors component. A large part of this literature is dominated by analysis done for the US economy (Baily, Bartelsman and Haltiwanger, 2001, Foster, Haltiwanger and Krizan, 1998 and Foster, Grim and Haltiwanger, 2013) and finds that although

labour reallocation plays an important role, within-firm dynamics dominate the cyclical behaviour of productivity within particular sectors. However most of this literature tends to focus on narrowly defined sectors (manufacturing being one of the most studied sectors) due to the availability of firm-level data. Therefore, the whole-economy implications are not always as clear cut.

Disney *et al* (2003) is perhaps the most authoritative work that discusses drivers of UK manufacturing TFP growth during the 1980s and 1990s recessions. They find that external restructuring, namely entry, exit and changing market share could explain around 50% of UK labour productivity growth within this sector between 1980 and 1992. We expand on this work by looking at the wider economy² up to 2011 and find that reallocation across firms could explain approximately 48% of labour productivity growth prior to 2007. However, after 2007 the contribution from labour reallocation declines and becomes negligible after 2009. This latter result is consistent with Broadbent (2012)'s assessment that resource reallocation has been slow to take place after the recent financial crisis. It is also in line with the literature that finds that financial crisis impede the movement of factors of production – Basu *et al* (2006) find for example that an impaired banking sector in Japan prevented capital from moving to the most productive companies and is likely to have affected aggregate TFP growth. Caballero and Hammour (2005) also find that crises freeze the restructuring process and that this is associated with the tight financial-market conditions that follow. They stress that the productivity cost of recessions adds to the traditional costs of resource under-utilization.

Impaired resource allocation following the financial crisis suggests that factors of production may not be in the most efficient places, which may limit the ability of the economy to take advantage of the changing composition of demand. However, as and when access to credit and the financial sector more generally recovers, one might expect the resource reallocation process to gather pace generating future productivity improvements.

See Appendix A for further details on this literature.

3 Data description

The analysis in the first part of the paper uses what we call the **sample dataset**. This is based on the ONS Annual Business Survey (ABS) since 2008, and the ONS Annual Business Inquiry (ABI) between 2002 and 2007. The ABS is an annual survey covering around two-thirds of the economy in terms of GVA and is composed of a questionnaire of around 60,000 businesses. It is principally used by the ONS to benchmark official GVA estimates. It replaced the very similar ABI in 2008.

Each survey includes information on Gross Value Added (GVA) as well as a range of other variables such as wages and capital expenditure. To construct real GVA for each firm we

² We use only Private Non-Financial Corporations (PNFCs), excluding the agriculture, energy, real estate and the public industries which seem to display volatile behaviour during the period covered in our sample. Also, Disney *et al* (2003) use establishment data from the Annual Respondents Database (ARD) while we use reporting unit data in our analysis. Harris and Moffat (2013) use a similar approach to Disney *et al* (2003) to discuss the role of foreign ownership on productivity growth.

merged in deflators provided by the ONS at the 2 digit Standard Industry Classification (SIC). The ABI also includes information on the number of employees up to 2007. Therefore for the post 2007 years we merge the ABS with the employment data from the ONS Business Register and Employment Survey (BRES)³. This involved matching the two datasets using the individual firm identifiers.

Our analysis covers the period between 2002 and 2011 and concentrates on Private Non-Financial Corporations (PNFCs), and excludes the agriculture, real estate and energy industries⁴.

To ensure that our analysis is representative of the UK economy as a whole we use employment and output weights provided by the ONS. Although these weights are a good approximation for the first part of our paper, they are inadequate for the second part. When considering the allocation of resources we need to understand the behaviour of newly formed firms and of those that fail. This information is missing in the ABS and ABI, but is included within the Inter-Departmental Business Register (IDBR). This is a live database of firms based on VAT and PAYE registrations.

The IDBR also contains employment and turnover data for the population of firms registered with it, but lacks information on GVA. With this in mind we construct a **population dataset** that exploits information on when a firm is created and when it dies using the population of registered firms from the IDBR. We have then matched in data from the *sample dataset*, and used the information on turnover, employment, size, age and sector already in the IDBR to impute estimates of GVA for the remaining non-sampled part of the population. Further information on how we construct these weights is provided in Appendix D.

A comparison of the sample and population datasets is presented in Table 1.

³ Both the ABS and BRES use stratified sampling techniques.

⁴ This is mainly because ABS and BRES do not provide full coverage of the agriculture, financial, real estate, energy and public sectors. Before 2004 these sectors are also unrealistically too volatile in our sample.

Table 1: Key features of datasets used

	Sample	Population
Number of companies of which:	39,000 to 50,000 per year	1,621,000 to 1,809,000 per year
• Small	26,000 to 35,000 per year	1,585,000 to 1,776,000 per year
• Medium	4,900 to 9,700 per year	26,000 to 31,000 per year
• Large	4,400 to 6,000 per year	5,900 to 6,600 per year
Time periods	2002-2011	2002-2011
Measure of labour productivity	GVA per employee	GVA per employee
Information sources	Annual Business Survey (ABS) and Business Register and Employment Survey (BRES) 2008-2011; Annual Business Inquiry (ABI) 2002-2007.	Inter-Departmental Business Register (IDBR), with imputed GVA estimates based on sampled values from the Annual Business Survey (ABS) and Business Register and Employment Survey (BRES) 2008-2011; Annual Business Inquiry (ABI) 2002-2007.
Variables included	Gross value added (GVA) Employment (year average) Total labour costs Capital expenditure	<i>Imputed</i> Gross value added (GVA) Employment (year average) Turnover

We use the same industries in the population dataset as we did in the sample datasets. For both data sets we consider outliers to be below the 1st or above the 99th percentiles of the labour productivity distribution.

3.1 Aggregate comparisons

As mentioned above, the ONS provide a firm-level weighting system with the ABS and ABI with which to scale up the sample estimates of GVA or employment to the aggregate economy. However, these aggregate estimates will not match official ONS National Statistics, since official estimates undergo a balancing process using a range of other data sources.

Charts 2 to 5 below show aggregate GVA growth, employment growth, labour productivity growth and the level of labour productivity estimated from the weighted sample dataset (blue lines). These are compared against two series. The first are published estimates by the ONS in their ABI and ABS publications (purple lines). The second are official ONS estimates of GVA and Workforce Jobs (WfJ, yellow lines). In both these series we use the same set of industries as included within our sample dataset.

The sample estimates appear to match employment estimates from the published ABS, ABI and WfJ relatively well (Chart 3). On the other hand, average growth in GVA in the sample dataset appears significantly stronger than the published ONS measures, although the broad trends are similar. A large part of this will reflect the relatively basic data cleaning process that we have

undertaken, relative to the steps that the ONS will take to produce their aggregate statistics (Chart 2). As a result, productivity growth in the sample is stronger over the full sample period compared to both published sources (Chart 5). The level of productivity, both in the sample and published ABI/ABS, appears to peak in 2008 whereas published ONS estimates of output per job peak in 2007 (Chart 4). Part of the difference may be due to the structural break that was introduced when the ONS moved from the ABI to the ABS in 2008.

The aim of this paper is to understand the labour productivity dynamics across firms, and not to produce aggregate level statistics. Therefore, given the broad trends are similar, we are confident that we are able to make inferences across firms based on the sample dataset that we have compiled.

Chart 2: Estimates of Gross Value Added (GVA) growth

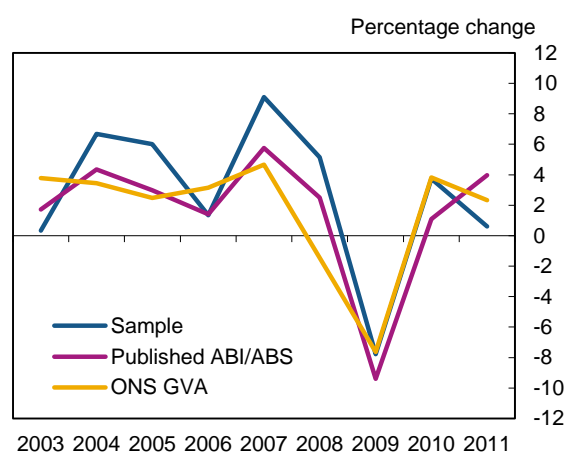


Chart 3: Estimates of employment growth

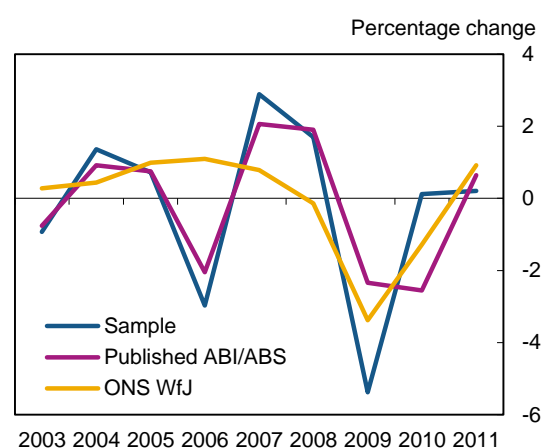


Chart 4: Estimates of labour productivity

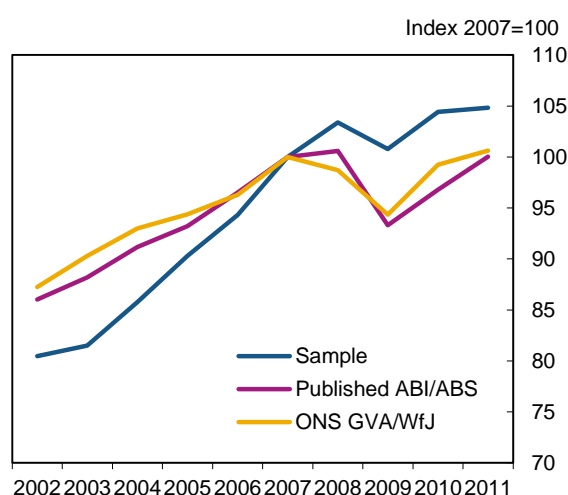
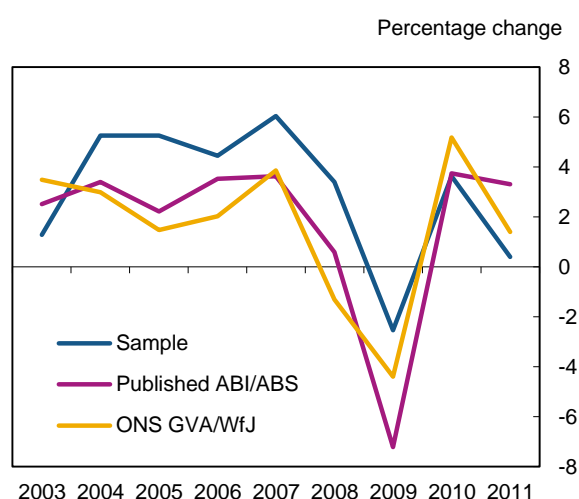


Chart 5: Estimates of labour productivity growth



Source: Bank calculations based on the sample dataset; ONS Annual Business Survey (2012); ONS estimates for Gross Value Added (GVA) and Workforce Jobs employment (WfJ). All estimates are for Private Non-Financial Corporations (PNFCs) excluding the agriculture, energy and real estate sectors.

4 Firm heterogeneity

This section discusses the degree of firm level heterogeneity and how it relates to aggregate productivity. First we look at the distribution of productivity across firms, sectors and sizes. Second, we look at the employment dynamics across different types of firms.

4.1 Distribution across firms

The distribution of labour productivity across firms is wide even before the recent financial crisis, as shown in Chart 6, and has shifted to the left since. Relative to 2007, 2009 saw a noticeable increase in a lower tail of firms reporting very low productivity levels. By 2011, a greater proportion of firms had lower productivity than they did in 2007.

The average gap between the top 5% of firms and bottom 5% of firms has been increasing over time, as shown in Chart 7. Between 2002 and 2006 this averaged at around £94,000 per head. By 2011 this had increased to £119,000 per head, an increase of 26%.

Chart 6: Kernel density of productivity distribution of firms, 2007, 2009 and 2011

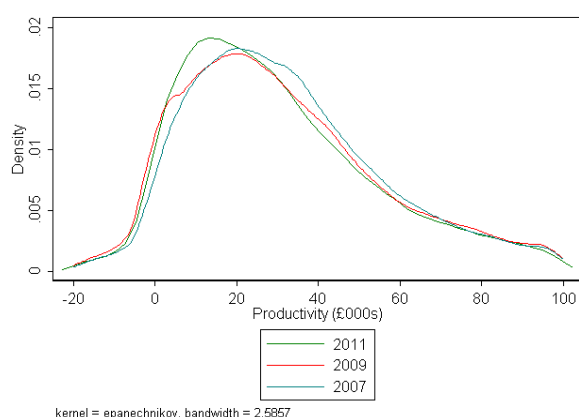
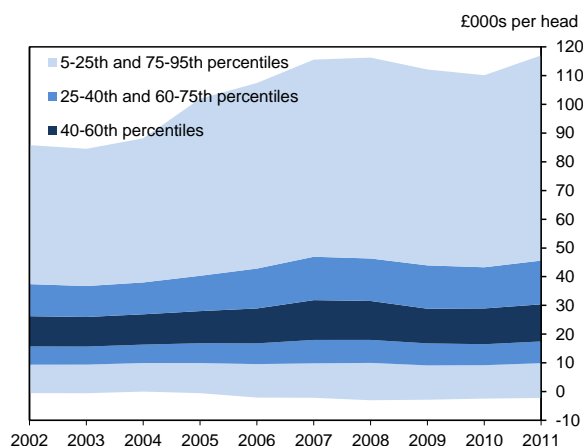


Chart 7: Productivity distribution of firms, 5th to 95th percentiles over time



After falling briefly during the worst part of the recession – reflecting a broad based hit to productivity, firms at the top of distribution have continued to increase their productivity levels. But measured productivity of firms at the bottom of the distribution has also gotten worse, with a higher proportion of firms reporting negative levels of productivity. Labour productivity for the central mass of firms is just below its pre-crisis level.

The distribution across firms within each industrial sector is also wide, and has increased in most cases since 2007. These are shown in Appendix B.

4.2 Distribution of loss-making firms

As shown in Chart 6, some firms reported negative labour productivity levels. These are firms that reported negative profits, measured as operating surplus here. Since Gross Value Added can be thought of as the sum of a firm's operating surplus and wage bill, this means that negative

profits have more than offset that firm's wage bill. This may occur when firms receive lower than expected turnover, which is unable to offset their non-labour costs.

Chart 8 shows the proportion of firms reporting negative profits over time. As shown, there has been an increase since 2007. This is broadly consistent with estimates based on financial accounts, reported in the August 2013 *Inflation Report*⁵. Appendix B looks at these developments by industry; Accommodation & Food and Construction saw some of the largest increases in the proportion of loss-making firms by 2009.

Chart 8: Percentage of firms with negative operating surplus

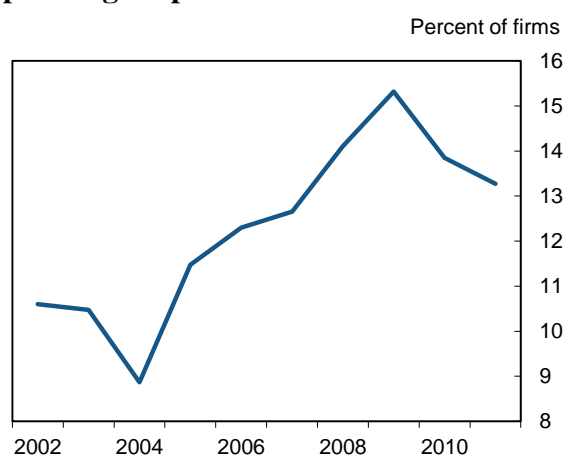
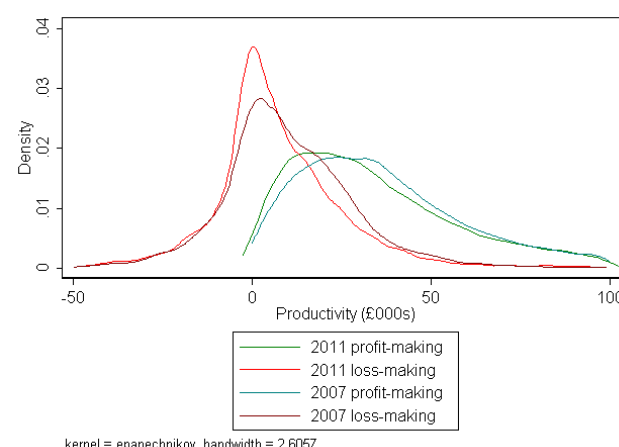


Chart 9: Productivity levels of firms making a profit and making a loss



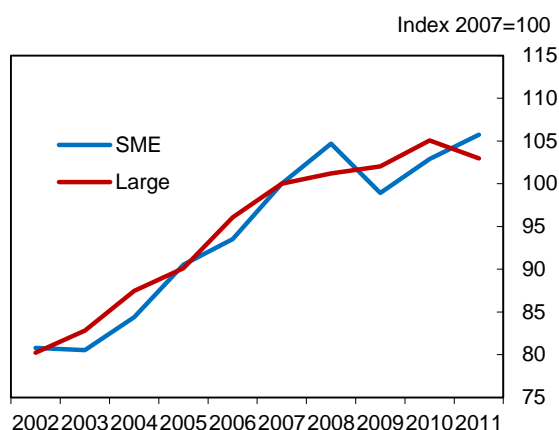
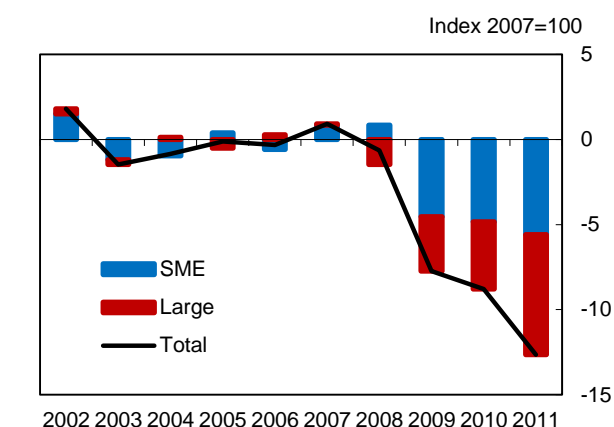
Firms making negative profits also had significantly lower productivity levels. Furthermore, between 2007 and 2011 not only did the proportion of firms making a loss increase, but average productivity of those firms also decreased. Chart 9 shows that the productivity distribution of loss-making firms moved towards the left after the crisis.

This is perhaps not surprising given that the UK has experienced lower firm death rates this recession compared to previous recessions. The implication being that fewer loss-making firms relative to past recessions will have gone bankrupt and exited the distribution.

4.3 Productivity across different firm sizes

The productivity experience across different firm sizes has also varied. The November 2012 *Inflation Report* used Company Accounts data to look at the productivity across firm sizes, and found that productivity across SMEs had fallen by more since 2008 than that of large firms. Crawford, Jin and Simpson (2013) found a similar result, and concluded that productivity growth, controlling for industry characteristics, had slowed by 7% for SMEs compared to no fall for large firms between 2007 and 2009 relative to their pre-recession trend.

⁵ It is important to note that the measure we use here is based on a firm's operating surplus, calculated as GVA minus total labour costs, and is different to accounting profit as measured in financial accounts.

Chart 10: Productivity by firm size**Chart 11: Contributions to the weakness in productivity relative to its pre-crisis trend**

Notes: Pre-crisis trend calculated as the linear trend between 2002 and 2007.

Chart 10 below shows the productivity level for SMEs, defined as firms with less than 250 employees, and large firms. As shown, productivity fell more sharply for SMEs at the onset of the crisis than for large firms, as the latter cut more on employment particularly in 2009. However, productivity growth of large firms has slowed considerably since then. By 2011 the weakness in productivity relative to its pre-crisis trend growth rate appears relatively evenly distributed across SMEs and large firms (Chart 11). The productivity distributions across firms within each of these categories are again wide, and have again increased significantly since 2007.

4.4 Employment behaviour

To investigate firms' employment behaviour further, this section looks at the output and employment patterns across different types of firms during and after the 2007/8 financial crisis.

For this exercise we divide the sample dataset into buckets based on their employment and output behaviour.⁶ Tables 2.1 to 2.4 show the percentage of firms in each employment and GVA growth category. Firms are weighted by the size of employment in the previous year. A firm is considered to have flat employment or flat GVA growth if their annual growth rate is between -5 and +5 per cent⁷. And red shading highlights categories where the percentage of firms is greater than 15%, and orange more than 7.5%.

Unsurprisingly we find that in 2009, following the onset of the financial crisis, there was a significant move away from firms that were growing and hiring, towards firms that were shrinking and firing. However by 2011 the proportion of firms firing workers had fallen significantly relative to 2009. In fact, a large concentration of firms emerged that had shrinking output but no change in employment. In 2011 the proportion of firms with shrinking output and flat employment was 22%, compared to an average of 11% between 2005 and 2007.

⁶ The sample of firms here is restricted to those who were included in the ABI/ABS survey in consecutive years, in order to calculate growth rates.

⁷ Although this range may appear large when considering aggregate changes in employment, at the firm level the variation in employment is much larger (as shown in Chart 12). The broad pattern shown in Tables 2.1 to 2.4 does not change significantly when these cut-off points are varied.

Table 2.1: Percentage of firms in each employment and output category, 2005-7

		Employment growth		
		Firing	Flat	Hiring
Output growth	Shrinking	14	11	7
	Flat	7	11	5
	Growing	12	19	16

Table 2.2: Percentage of firms in each employment and output category, 2009

		Employment growth		
		Firing	Flat	Hiring
Output growth	Shrinking	20	18	4
	Flat	6	12	2
	Growing	10	18	8

Table 2.3: Percentage of firms in each employment and output category, 2010

		Employment growth		
		Firing	Flat	Hiring
Output growth	Shrinking	10	20	6
	Flat	6	14	5
	Growing	8	19	13

Table 2.4: Percentage of firms in each employment and output category, 2011

		Employment growth		
		Firing	Flat	Hiring
Output growth	Shrinking	11	22	6
	Flat	4	19	5
	Growing	6	17	10

To understand this change in behaviour, Chart 12 plots the distribution of firms and their associated level of employment and GVA growth. Each bubble represents a firm, but its size is determined by total employment in the previous period and weighted by how representative that firm is in the population (using the ONS weights). As shown the distribution of firms is large, however does change between the pre- and post-crisis periods. In fact, the null hypothesis that the distributions of GVA, employment and productivity growth are equal in both periods is rejected at the 1% significance level using a Kolmogorov-Smirnov test. This chart also includes lines of best fit for each period, although they are hard to see. Chart 13 shows these same lines of best fit but without the bubbles and on a smaller axis. As shown, the relationship between GVA growth and employment growth appears to flatten after the crisis. This suggests, that for a given level of GVA growth, employers were both less likely to hire and less likely to fire workers. Since a greater proportion of firms had negative GVA growth post crisis, this change in behaviour is likely to have contributed significantly to the weakness in productivity.

Chart 12: Distribution of GVA and employment growth outcomes

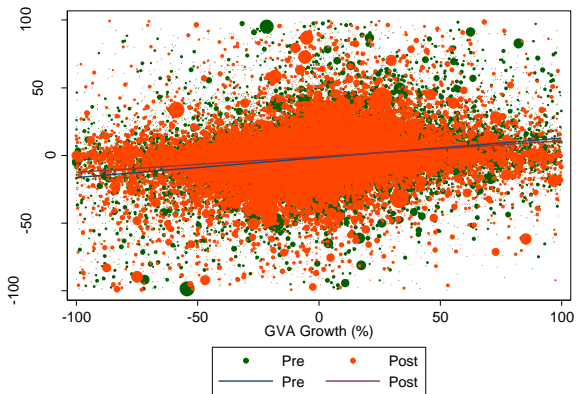
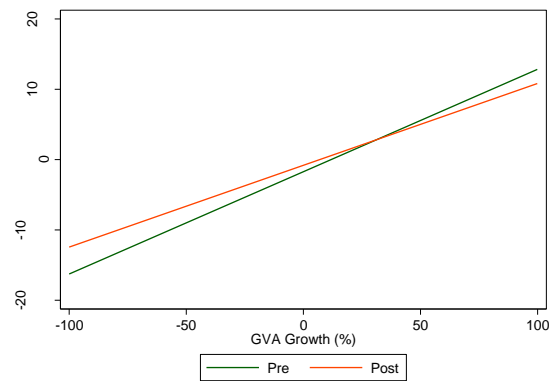


Chart 13: Linear relationship between GVA and employment growth outcomes



Notes: ‘Pre’ refers to the period 2005 to 2007. ‘Post’ refers to 2008 to 2011.

Tables 3.1 and 3.2 provide an estimate of the contribution to annual productivity growth of each of the categories described above for 2005-7 and 2011 – to help understand the role of each group in recent productivity outturns. These are calculated by multiplying the GVA shares of each group with the average change in productivity in each group. Cells with positive productivity growth are highlighted in green, and negative growth red. The total of the contributions shown in these tables will not sum to the aggregate growth rate of productivity, since they do not include any new entrants or firms that die in a particular period or firms that are included in the ABI/ABS sample for one year only. The key finding from this is that the largest drag on annual productivity growth relative to 2005-7 came from the group of firms with shrinking output and flat employment. This contributed -7pp to annual growth in productivity in 2011 compared to -3pp on average between 2005 and 2007. This suggests that firms that have been holding on to labour in the face of weak demand have been a large driver of the weakness in productivity in 2011, rather than firms actively hiring.

Table 3.1: Contribution to productivity growth, 2005-7

		Employment growth		
		Firing	Flat	Hiring
Output growth	Shrinking	-1	-3	-4
	Flat	1	0	-1
	Growing	6	5	2

Table 3.2: Contribution to productivity growth, 2011

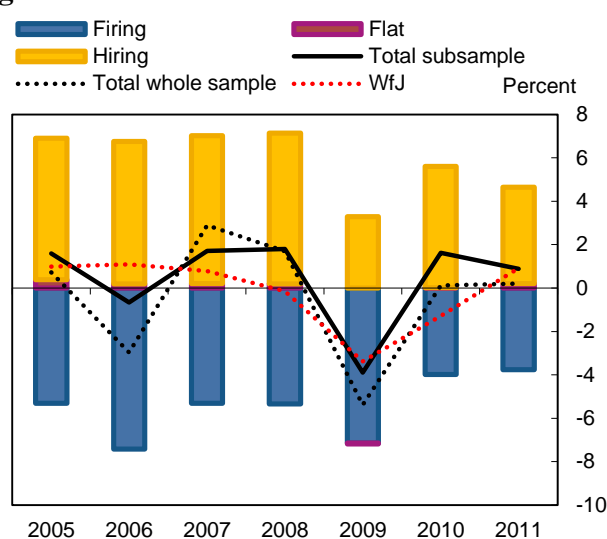
		Employment growth		
		Firing	Flat	Hiring
Output growth	Shrinking	-1	-7	-3
	Flat	1	0	-1
	Growing	3	6	1

Chart 14 shows the contribution to net employment growth from firms defined as ‘firing’, ‘flat’⁸ and ‘hiring’. These categories will make up a subset of all the firms in our sample, since in order

⁸ As described earlier, a firm is considered to have flat employment or flat GVA growth if their annual growth rate is between -5 and +5 per cent.

to calculate growth rates you need to observe a firm for two consecutive periods. And there are many firms that are either only sampled in one particular year or drop out of the sample at a particular point in time. Therefore, Chart 14 compares the total from this subsample with the aggregate change across the whole sample (dotted black line) and the ONS Workforce Jobs (WfJ) measure (dotted red line). On average before the crisis net firing and net hiring by firms were relatively substantial, contributing around -6% and +7% to net employment growth respectively. These flows are significantly larger than the total net change. But this is also true when looking at gross flows into and out of employment from the Labour Force Survey (LFS). Between 2005 and 2007 gross flows into and out of employment stood at around 7.5 million or 26% of total employment each year.

Chart 14: Contribution to total employment growth



As shown in Chart 14, the contribution from firms that are hiring fell from 2009 and increased substantially from firms that were firing. In 2011, firms that were hiring contributed around 4pp to total employment growth, higher than we would have expected in a period of weak output growth but still below a pre-crisis average of 6pp. However, this has largely been offset by a smaller negative contribution from firing; the contribution of firing to employment growth in 2011 was less negative by around 2pp compared to the pre-crisis period. We discuss the impact of labour reallocation - through firing and hiring - on productivity growth in the next section of this paper.

A key finding from this analysis is that the largest drag on annual productivity growth relative to 2005-7 came from the group of firms with shrinking output and flat employment. This suggests that firms that have been holding on to labour in the face of weak demand have been a large driver of the weakness in productivity up to 2011.

One of the difficulties with putting too much weight on this interpretation is that aggregate employment has been growing strongly in 2012 and 2013 despite relatively weak output growth. Unfortunately, we only have access to firm-level data up to 2011 so are unable to shed light on more recent developments.

Overall, there are two key findings from this section. First, the productivity distribution across firms has widened and shifted towards the left – which has also been matched by an increase in loss-making firms over the recession. While firms of different sizes have adjusted at different speeds, the weakness of productivity relative to its pre-crisis trend growth rate appears evenly distributed across SMEs and large firms. Second, analysing firms' employment behaviour suggests that firms holding on to labour, despite falling output, have been a large driver of the weakness in productivity up to 2011.

In the next section, we consider the contribution to aggregate productivity growth from the dynamic reallocation of labour and capital cross firms.

5 The effect of firm-level restructuring and creative destruction

It is widely accepted that productivity growth in the long run is driven by restructuring and a reallocation of the factors of production. The simple intuition behind this is that reallocation allows factors of production to move from low to high-productivity firms. Indeed this has been demonstrated by a large and growing literature for a number of countries and using a variety of methods. One form of reallocation has often been termed ‘creative destruction’ – the process where less productive firms go bankrupt and release resources to more productive firms to grow and expand.⁹

There are also important reasons why the process of reallocation could be impeded after a financial crisis. For example, Broadbent (2012) argued that a sclerotic banking sector that restricts credit availability could impede the allocation of resources from one sector to another. And when faced with an asymmetric demand shock induced by the financial crisis, the need for reallocation may have also increased at the same time. Therefore, quantifying the effects of reallocation on labour productivity, and tracing these through time, must be important to our understanding of productivity following the most recent financial crisis.

5.1 Methodology

To quantify the effects of reallocation and creative destruction, we follow the method of decomposing changes in aggregate labour productivity developed by Baily *et al* (2001)¹⁰ – hereafter referred to as BBH. Under their methodology, changes in aggregate productivity can be decomposed into four parts: productivity changes within individual firms, movement of labour from one firm to another, and movements into new firms and away from dying ones. The advantage of this methodology is that it is tractable and easy to interpret, since the four terms analytically sum up to aggregate productivity changes on the left-hand side.

The formula that is applied to our analysis is shown below:

⁹ See Caballero (2007) for a comprehensive literature survey.

¹⁰ Baily *et al* (2001) is only one of many related methods of decomposition; we chose this method mainly because of its robustness, simplicity and tractability. For further information on alternative methods of decomposition, see Foster *et al* (1998).

$$\begin{aligned}
\Delta \Pi_t = & \sum_{i \in Surv} \bar{\phi}_i \Delta \pi_{i,t} && \text{..... within firm} \\
& + \sum_{i \in Surv} \Delta \phi_{i,t} (\bar{\pi}_t - \bar{\Pi}) && \text{..... reallocation between surviving firms} \\
& + \sum_{i \in Entry} \phi_{i,t} (\pi_{i,t} - \bar{\Pi}) && \text{..... reallocation to new firms} \\
& - \sum_{i \in Exit} \phi_{i,t-1} (\pi_{i,t-1} - \bar{\Pi}) && \text{..... reallocation from exiting firms}
\end{aligned}$$

where $\pi_{i,t}$ is GVA per head in firm i at time t ,

Π_t is aggregate GVA per head at time t ,

$\phi_{i,t}$ is the employment share of firm i at time t , and

a bar over a variable indicates the average of the variables across $t-1$ and t .

The first term of this decomposition shows the effect of productivity changes within individual firms on aggregate productivity. Changes in individual firms' productivity are weighted up using each firm's employment shares, which means that changes within large firms will have a bigger impact on the aggregate than changes within small firms.

The next three terms together capture the effect of labour reallocation, albeit from different sources. The second term reflects how changes in labour shares of existing firms affect the aggregate (a proxy for changes in market share). It is positive when the labour share of a firm increases (decreases) and that firm is more (less) productive than the average firm in the economy. The third and fourth terms refer to the contribution of births and deaths. The intuition here is more straight forward: it calculates the impact on aggregate productivity of a firm birth or death by comparing their productivity level with that of the average firm, and weights this by how much labour they employ. For instance, when a new firm is more productive than the average firm, aggregate productivity increases; but when a dying firm is more productive than the average firm, aggregate productivity decreases because their workers are likely to be moving to less productive firms on average.

5.2 Data requirements and description

An important prerequisite of applying the BBH methodology is to have output (i.e. GVA) and employment information for *each and every* firm in the economy. This is a major challenge for most UK firm-level datasets. However, our *population dataset*, described in Section 2, overcomes this shortfall by combining the sample information on GVA and employment from the ONS Annual Business Survey and Annual Business Inquiry with the longitudinal characteristics of the population of registered firms in the Inter-Departmental Business Register

(IDBR) dataset. The population dataset therefore allows us to identify when firms are born or die.

In this paper we impute non-sampled firms' GVA from the IDBR to generate the BBH decomposition. Our imputation methodology is described below.

The first step is to calculate the average GVA to turnover ratio for well-defined groups of sampled firms. A key difference between the existing literature and our paper is that we use groups of firms which are defined by their 4-digit SIC industry code, their size and age. Given that this decomposition is aiming to analyse births and deaths it is important to impute new firms' GVA using information obtained from similar aged firms.¹¹ We then apply this ratio to non-sampled firms within the same group - this allows us to impute GVA by multiplying this ratio to non-sampled firms' observed turnover from the IDBR.¹² Further details of our imputation methodology can be found in Appendix D.

Data imputation is not ideal and has its caveats. White, Reiter and Petrin (2012) discuss some of the drawbacks and benefits of data imputation for the U.S. Census Manufacturing data. Ideally one would work with observed firm-level GVA or the capital stock of all firms in the economy. But as this is not possible previous research (Harris, 2002 and Disney *et al*, 2003) has used available data for similar industry and size companies to construct un-sampled variables such as local-unit capital stock or GVA. Criscuolo, Haskel and Martin (2003) also describe how the ONS sometimes imputes reporting units' information which is not surveyed with similar industry-region-size data (this imputation process is called expansion).

As a crosscheck on our exercise we compare our estimated proportion of births and deaths with those published by the ONS. In line with results obtained from ONS's own calculations, we find that the proportion of dying firms has increased in this recession, compared to the average of the last 10 years (black dotted line in Chart 15). But this increase is much more muted than that in the 1990s and the 1980s, as indicated by solid grey line based on aggregate data in the same chart. Firm entry has also decreased and, despite ticking-up in 2011, has remained low since 2007.

We also find that on average the productivity of firms that die is lower than that of existing firms and new firms (Chart 16); this gap has narrowed markedly during the financial crisis. Interestingly the level of productivity of new firms that survive for longer than one year has been similar to that of survivors¹³ since 2005. On average every year, 98% of these companies have less than 50 employees and therefore account for a small share of aggregate employment.

¹¹ See Fort, Haltiwanger, Jarmin and Miranda (2013) for a discussion of the importance of size and age.

¹² We conducted several sensitivity analysis by using different size, age and industry groupings. We also compared our results using imputed data with a decomposition obtained using the weighted sample. The results are broadly similar with the main conclusions remaining unchanged.

¹³ Perhaps this is not surprising as we deflate newly born firm's GVA with a 4 digit industry deflator which US research proved to be higher than the price with which the new firms enter the market, to try to undercut competitors.

Chart 15: Proportion of firms that exited and entered the market

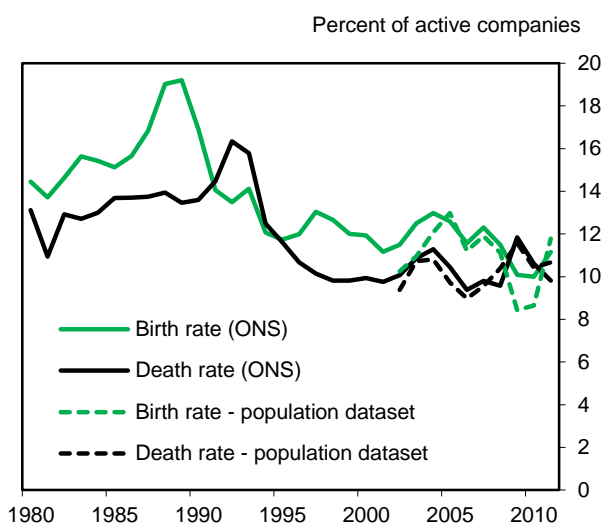
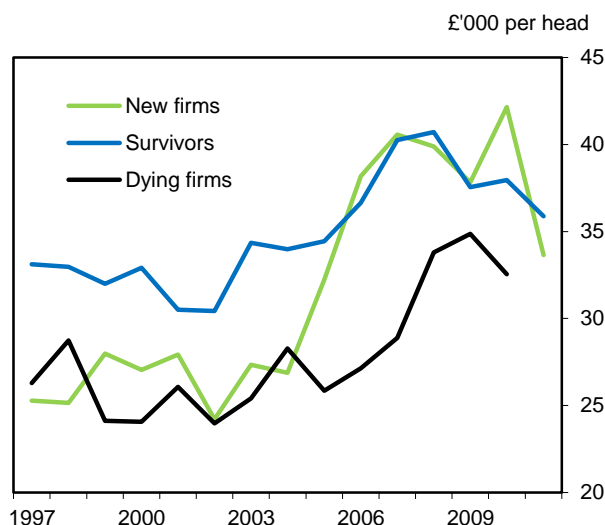


Chart 16: Productivity level by age



Source: ONS Business Demography data and Bank calculations; estimates based on the population dataset.

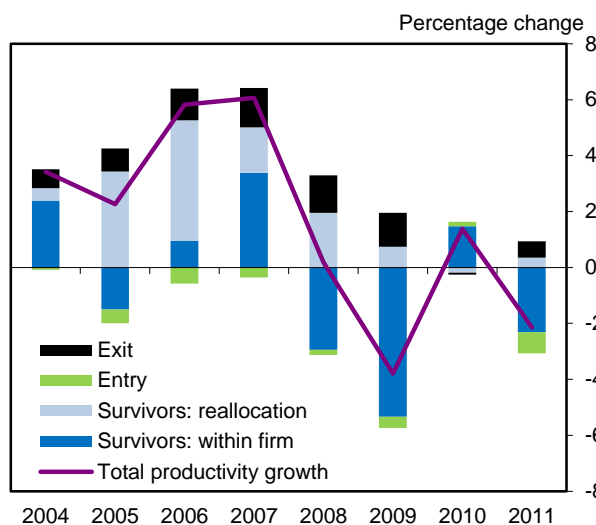
5.3 Baseline results

The result of the BHH decomposition is shown in Chart 17 below. Among the four components, the within-firm productivity of surviving firms (shown as dark blue bars here) is the most important in driving the cyclical pattern of aggregate productivity growth; this is consistent with findings by both BBH for the US and Disney et al (2003) for the 1990s in the UK.

The contribution of labour reallocation among survivors (shown as light blue bars here) has been positive and relatively stable between 2002 and 2009. But after 2009, the contribution from reallocation has become negligible. One interpretation of this is that the recession has hampered reallocation of labour across the economy.

The reallocation of labour towards new firms does not contribute positively to year on year productivity growth; in fact, for most years it is negative. This is likely to be because start-ups can take some time to become profitable due to the fixed costs involved in entering the market place. It is also worth noting that our estimates are based on reporting units, rather than at the plant level, which means the effect of any internal restructuring (from closing and opening plants) is likely to be captured in our 'within-firm' component.

Chart 17: Decomposition of annual labour productivity growth



Dying firms, on the other hand, do contribute positively to aggregate productivity, since on average they have a lower level of productivity relative to the average.

Another way to summarise the information in Chart 17 is to cumulate the contributions over two distinct periods: before and after the financial crisis.

In Chart 18, we have defined the pre-crisis period as 2002 to 2007 and the post-crisis period as 2008 to 2011¹⁴. We then calculate the average contribution over that period.

This allows us to summarise the contributions to aggregate productivity growth over these periods and gives new firms time to build up their productive capacity. However, these periods are still relatively short. In order to understand the long-term contribution of net entry to productivity growth it may be more appropriate to look over longer time horizons.

Over the pre-crisis period, productivity improvements within firms contributed around 2.2pp to aggregate productivity growth, while reallocation of labour across firms contributed around 1.2pp. The contribution from net entry of firms was also important, adding 1.7pp in total to aggregate productivity growth, most of which was due to unproductive firms exiting the economy.

After the crisis, this split changed significantly. Within-firm changes in productivity were the most significant component, subtracting 1.6pp from aggregate productivity growth. Reallocation of labour still contributed positively to growth, though it halved in magnitude to 0.5pp. And entry and exit of firms did little to boost lost productivity growth. Deaths of unproductive firms added 0.4pp to aggregate productivity growth, and new entrants subtracted a roughly equivalent amount from it. This might reflect the fact that there were fewer firms born after the crisis (middle row of Table 4.1) and their productivity level was, on average, 11% lower than aggregate productivity (middle row of Table 4.2). The relative productivity gap between surviving and dying firms had narrowed: while during the 2002-2007 period the average dying firm had a productivity level 33% lower than aggregate productivity, after the crisis this was only 18% lower (bottom row in Table B).

¹⁴ For each period, we compare the ending year's productivity levels and labour allocation across firms with those of the beginning year. In other words, for the Great Stability period, 't' in our formulae would be 2007, while 't-1' would be 2002; for the Great Recession period, 't' would be 2011 and 't-1' 2007.

Chart 18: Decomposition of productivity growth pre and post crisis

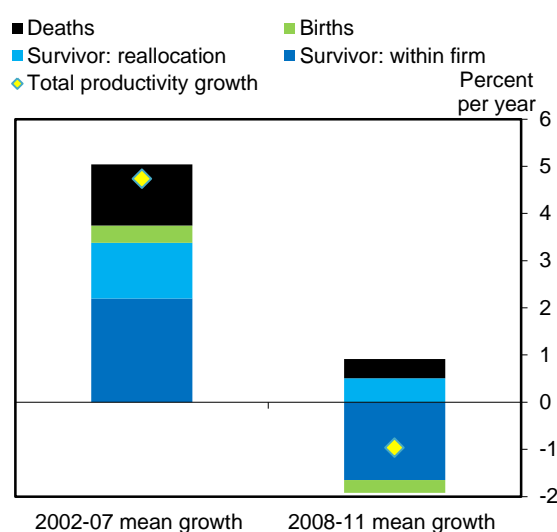


Table 4.1: Average number of firms per year

	2002-07		2008-11	
	000s	% total	000s	% total
Survived	934	75%	1,123	77%
Born	169	14%	165	11%
Died	136	11%	169	12%

Table 4.2: Average productivity of firms

	2002-07		2008-11	
	£000 / year	% diff*	£000 / year	% diff*
Aggregate	38.1	--	41.6	--
Born	41.1	+8%	37.0	-11%
Died	25.3	-33%	34.2	-18%

* percentage differences in productivity relative to the aggregate benchmark

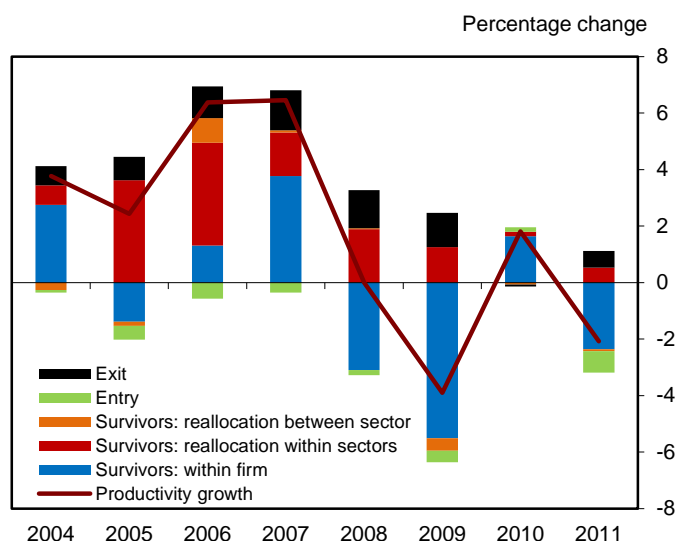
Overall, our decomposition shows that, of the 6% difference in aggregate productivity growth after the crisis, around two thirds can be explained by negative productivity growth within firms, while one third can be explained by less – or less effective – reallocation of labour, and slower ‘creative destruction’.

5.4 How much reallocation occurs across different sectors?

In the aftermath of the financial crisis of 2007/8, the UK experienced a series of asymmetric shocks affecting different sectors of the economy. For example, the 25% depreciation of Sterling would have helped exporting sectors like manufacturing, while harming import intensive ones like retail. Therefore, an interesting question is whether the pattern in reallocation we have seen is due to reallocation across different sectors, or simply reallocation within sectors.

Past literature has demonstrated that most reallocation effects tend to come from movements in labour between different firms *within the same industry*. Indeed this is what we find in our exercise too. Chart 19 plots the result of an evolved decomposition method, whereby we split the reallocation term into reallocation of labour between firms in the same sector (in red), and between different sectors (in purple). Appendix E goes through the methodology in detail.

Chart 19: Decomposition of productivity growth, including sector reallocation



The sectoral reallocation bars contribute very little to aggregate productivity growth. This does not necessarily mean that labour has failed to move across sectors – it merely means that any labour that did move did not result in large changes in aggregate productivity, perhaps because those movements were between sectors of similar levels of productivity, or that large movements to productive sectors were offset by movements into unproductive ones as well.

5.5 Putting our results into context

The above analysis suggests that up to a third of the productivity weakness can be explained by a slowing down in reallocation. But is this slowdown a typical occurrence during recessions as suggested in the literature? Without data covering previous recessions it is hard to understand how typical this behaviour has been.

In order to understand whether this recession is any different from the past we do two things. First we use data on manufacturing firms which goes back to the 1970s. Although this is not ideal as the manufacturing sector has not seen as protracted a period of weak productivity growth as the rest of the economy, it should be informative nevertheless. Second, we construct a counterfactual exercise to demonstrate what might have happened to labour productivity had the level of firm deaths picked up in a similar way to the 1990s recession.

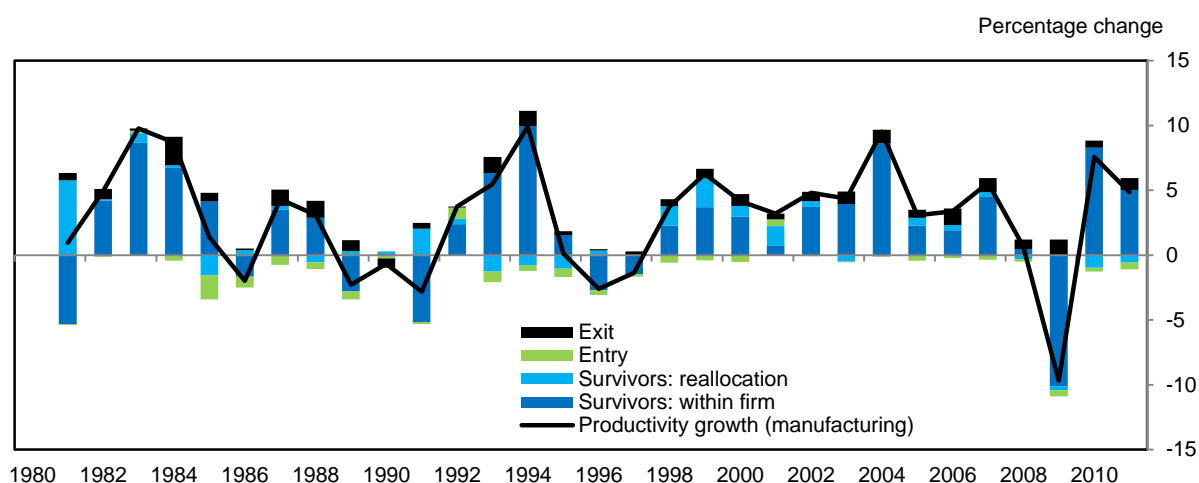
5.5.1 Manufacturing sector during three recessions

Chart 20 plots the same productivity decomposition as in the previous sub-section, but this time for manufacturing firms only between 1979 and 2011. Looking at past recessions, there are several points which stand out.

First, in line with the whole-economy results, the most cyclical component of productivity is changes within firms. However, it is important to note here that our dataset is at the firm, and not the plant, level. Existing studies have tended to use plant-level datasets when analysing the manufacturing industry, which may explain why our reallocation terms are much smaller than the within-firm term. Second, the reallocation component is strongest in 1981, 1991 and 1999-2001 which correspond with periods of recovery years for the UK economy as a whole. Third, the effect coming from companies dying has been positive for many years, reflecting the long, secular downsizing of UK manufacturing sectors over time. Although modest, the

manufacturing sector seems to have experienced a slightly lower level of labour reallocation in 2009 relative to 1980 and 1990.

Chart 20: Decomposition of annual labour productivity growth for manufacturing



5.5.2 The impact of higher firm survival on aggregate productivity

As discussed in the Bank of England August 2013 *Inflation Report*, the number of company failures has been unusually low since 2008 given the weakness in demand. There are several factors that may have helped firms survive over this period, for example the low level of Bank Rate, weak real wage growth and any forbearance by HMRC and the banks¹⁵. The extent and reasons why firms have been better able to survive this recession is beyond the scope of this paper. However we can conduct a simple experiment to illustrate what the impact higher corporate bankruptcies may have had on measured labour productivity growth.

To do this, we use our data set to imagine a counterfactual scenario which assumes that a larger number of companies would have failed after the financial crisis. To begin with, we first construct a simple OLS regression model of aggregate net entry (solid lines in Chart 21) using GDP growth and its lags between 1985 and 2007. We then use the estimated coefficients to project this relationship over the recent recession, to simulate what the net entry of firms would have looked like had the economy behaved as in the 1990s (dotted lines in Chart 21). And, using this simple model, we back out an expected death rate. This counterfactual death rate is around 5pp higher than that which we observed during this crisis.

¹⁵ See Arrowsmith *et al* (2013) for further details on a recent Bank of England investigation into bank forbearance in the SME sector.

Chart 21: Net entry and death rate counterfactual scenario

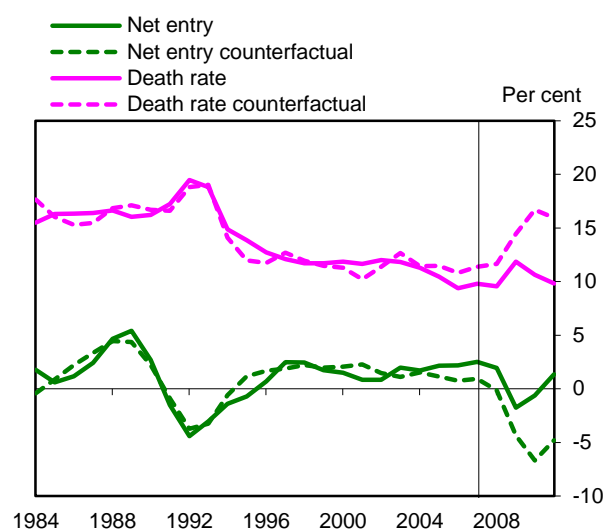
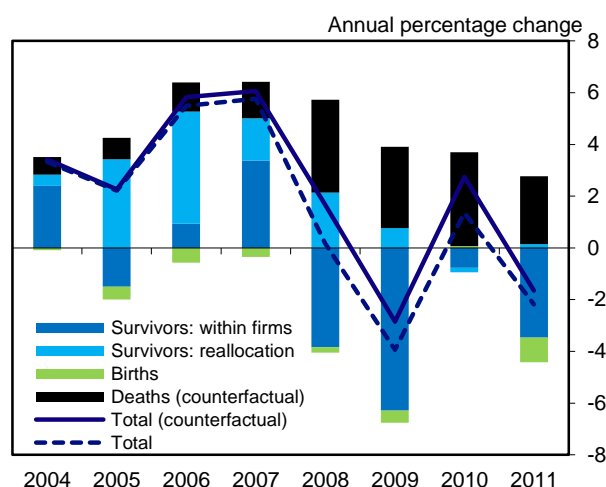


Chart 22: Decomposition of annual labour productivity growth: counterfactual



To generate the impact on aggregate productivity, we assume that additional firms would have died in our population dataset in order to increase the aggregate death rate by 5pp. There are a number of ways of doing this. One way is to assume that the most unproductive firms in the economy go bankrupt first. Under this assumption, the impact of this on the aggregate productivity level is shown in the blue solid line in Chart 22. This higher death rate has increased the size of the black bars from Chart 17 - pushing up on annual productivity growth. The cumulative difference between the dotted and the solid line is sizable at around 5% on the level of productivity by 2011.

While measured productivity is higher, this scenario is unlikely to have been optimal. Even if these additional firms arithmetically reduced aggregate *measured* productivity through a batting average effect, it is far from clear whether the weakness in productivity amongst these firms is temporary or more persistent. It may well be that we could see a strong recovery in productivity amongst this lower tail of firms as demand recovers. Therefore, higher firm survival may be a beneficial result of the economy's ability to better smooth through temporary demand conditions. In addition, there are likely to have been large welfare costs associated with a higher level of bankruptcies and any additional unemployment. On the other hand, to the extent that these firms are not viable in the longer term, their continued survival may impede any reallocation of resources to more productive uses. It is not possible to differentiate between these hypotheses with the data we have available. Instead the aim of this counterfactual experiment is to better understand the scale of the potential direct impact on measured productivity arising from the different behaviour of the UK economy relative to the 1990s.

6 Conclusion

There are a number of key findings that stand out from our work.

First, we find that a large part of the weakness in productivity can be accounted for by the fact that the proportion of firms with shrinking output and flat employment doubled from 11% in

2005-07 to 22% in 2011. This suggests that firms were able to respond flexibly to weak demand conditions by retaining staff at the expense of measured productivity, suggestive of an opening up of spare capacity within firms. This result implies a more *temporary* or *cyclical* explanation of the productivity slowdown in the UK, as these firms may have been well placed to increase production quickly without hiring as demand conditions started to improve. However, our analysis only extends to 2011, and the strength of recent hiring behaviour since 2012 implies that this may now be less of a factor.

The second key finding relates to the role resource allocation played in the slowdown of labour productivity growth. We find that labour reallocation and within-firm productivity improvements were equally important in driving productivity growth between 2002 and 2007. However since the crisis, the role of reallocation fell significantly while the contribution of within-firm productivity to aggregate productivity growth turned negative. Taken together, we find that a third of the shortfall in labour productivity since the crisis can be attributed to slower resource reallocation between firms. The rest is due to productivity weakness within firms. This result is indicative of a more *persistent* interpretation of the UK productivity slowdown. But since our data are at the firm and not plant level, these within-firm changes will not include any reallocation effects across individual plants that sit within particular firm entities and may, therefore, underestimate the scale of these reallocation effects. The speed at which labour productivity is able to grow in the short to medium term may be limited by the extent to which impaired allocation of resources across the economy continues to be a binding constraint. But exactly how companies and resources will respond as demand recovers remain key questions for the economic outlook.

Overall, this paper provides a comprehensive addition to the existing literature examining the UK ‘Productivity Puzzle’ and the role of resource allocation. However, there are a number of avenues that future work may aim to explore. The methodology for decomposing labour productivity into its drivers can be improved in several directions. For example, working with plant data rather than reporting units would be an important development. Panel econometric analysis could help further understand the drivers of productivity developments, although our preliminary explorations were difficult to interpret. In addition, in our paper we have worked with firm-level estimates of labour productivity, however examining what has happened to firm-level total factor productivity (TFP) by taking into account changes in capital per worker is also important.

Appendix A: Summary of the literature

Study	Country	Sample	Findings
Disney <i>et al</i> (2003)	UK	1980-1992- plant level data (ARD)	<p>They find that external restructuring accounts for 80-90% of establishment TFP growth – much of this effect comes from multi-establishment firms closing down poorly performing plants and opening new highly performing plants.</p> <p>All of the productivity gains among single establishments came from the entry of more efficient establishments and the exit of less efficient ones. Among multi-establishment firms, about half of productivity growth was due to differential productivity growth among survivors, whilst the other half was due to the closure of low productivity establishments and the opening of higher productivity ones. Thus much of the overall net entry.</p>
Harris and Moffat (2011)	UK	1997-2008- plant level data (ARD)	<p>Reallocation of output shares towards highly productive industries and the opening of highly productive plants explain most of productivity growth between 1997 and 2008. (Use ARD as source of data)</p>
Caballero R. and M. Hammour (2000)	US	Model calibrated to US manufacturing data	<p>Based on a combination of theory with empirical evidence on gross job flows and on financial and labour market rents, the authors find that, cumulatively, recessions result in reduced rather than increased restructuring, and that this is likely to be socially costly once we consider inefficiencies on both the creation and destruction margins. They find that the productivity costs adds to the traditional costs of resource under-utilization.</p>
Baily <i>et. al</i> (2001)	US	1972-1988- firm level data	<p>Firm-level productivity more pro-cyclical than aggregate productivity. Entry and exit does not have a large impact on productivity growth due to low employment shares.</p>
Levine O. And M. Warusawitharana (2014)	UK, US, Spain and Italy	1999-2009- firm level data (BvD)	<p>Find an economically important role of financial constraints on output during the most recent crisis. They claim that reduction in debt reduces investment in TFP increasing projects by firms thereby leading to lower TFP and labour inputs.</p>
Ben Broadbent (2012)	UK	1997-2012- industry level data	<p>Misallocation of capital across sectors against a background of disfunctioning financial sector is likely to explain the persistent weakness of aggregate productivity between 2009 and 2012.</p>

Appendix B: Productivity distributions by industrial sector

Chart B.1: Manufacturing (C)

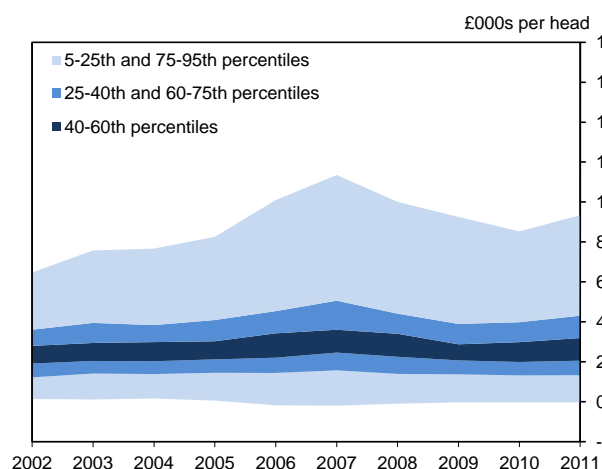


Chart B.2: Construction (F)

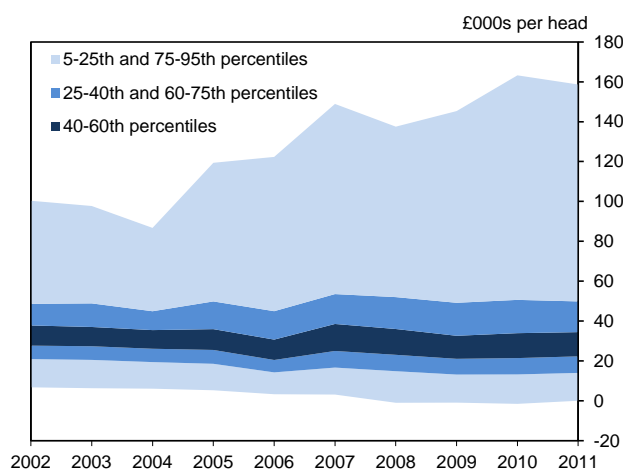


Chart B.3: Wholesale and Retail (G)

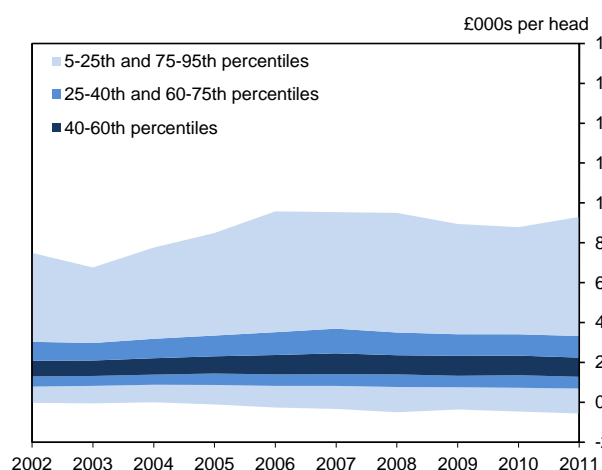


Chart B.4: Transport and Storage (H)

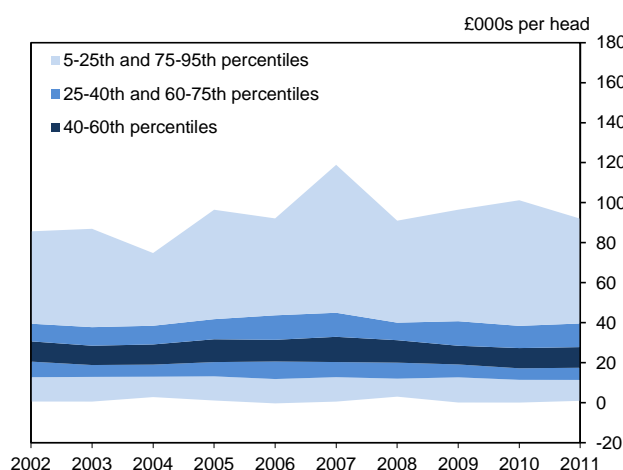


Chart B.5: Accommodation and Food (I)

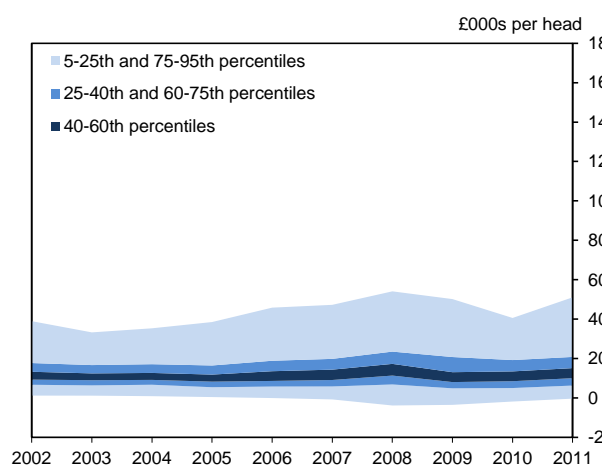


Chart B.6: Info and Communication (J)

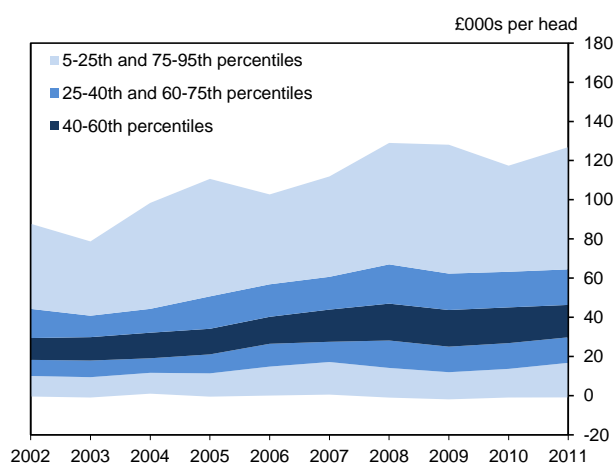


Chart B.7: Professional and scientific services; and admin services (M & N)

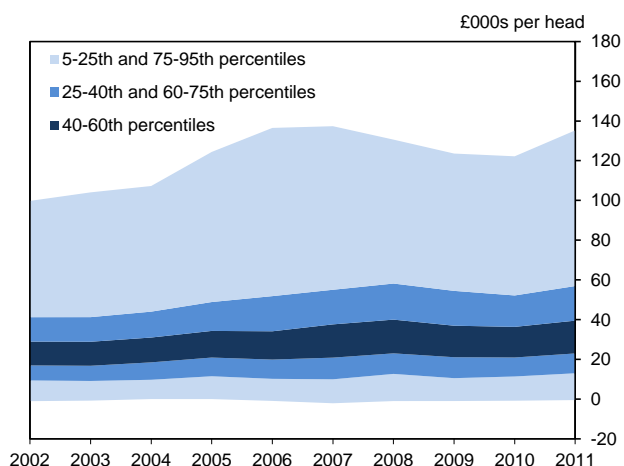


Chart B.8: Arts and Leisure Activities (R)

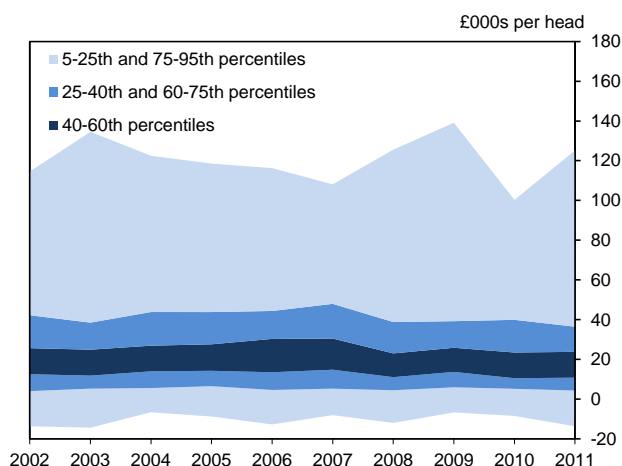


Chart B.9: Other services (S)

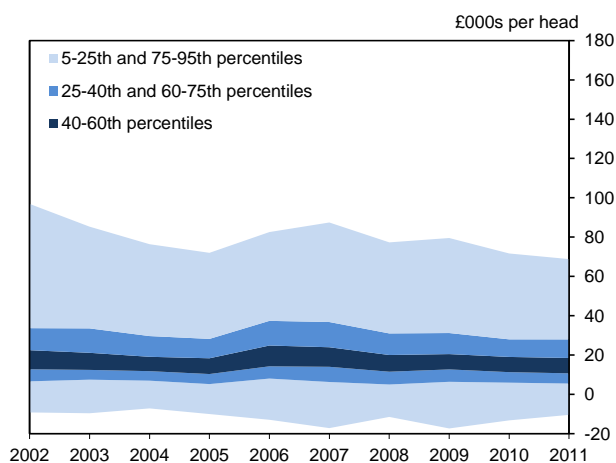
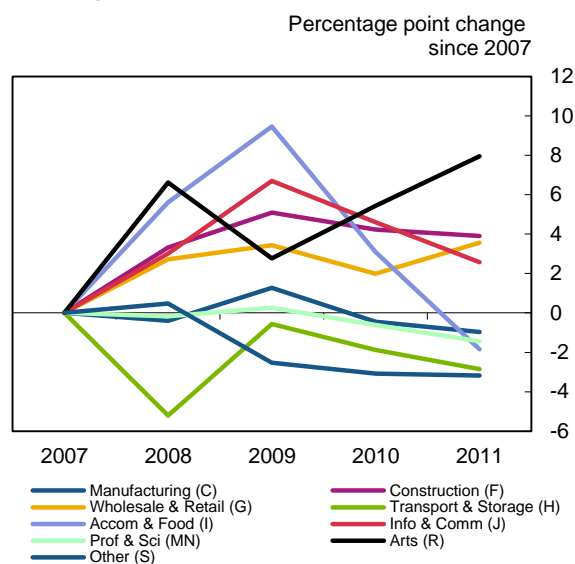


Chart B.10: Change in the percentage of firms with negative operating surplus by industry



Appendix D: How our imputation compares to ONS weights

The following provides further details on our imputation methodology and compares these with the ONS weighting methodology.

Assume the following data is available for a particular ‘bucket’ of firms, defined by their four-digit sector, size and age group.

Firm number	GVA	Turnover (NUL file)	Turnover (DAT file)
1	GVA_1	$T(nul)_1$	$T(dat)_1$
2	(N/A)	$T(nul)_2$	(N/A)
3...N-1	(N/A)	$T(nul)_3 \dots T(nul)_{N-1}$	(N/A)
N	GVA_N	$T(nul)_N$	$T(dat)_N$

To calculate aggregate GVA, the ONS – in essence – uses the following method:

$$\begin{aligned}
 &GVA_{TOTAL} \\
 &= (GVA_1 + GVA_N) \times \frac{\sum_{i=1 \dots N} T(nul)_i}{(T(nul)_1 + T(nul)_N)} \\
 &= (GVA_1 + GVA_N) \times \frac{N}{2} \times \frac{\text{Average of } T(nul) \text{ in population}}{\text{Average of } T(nul) \text{ in sampled firms}} \\
 &= (GVA_1 + GVA_N) \times aweight \times gweight_{turnover}
 \end{aligned}$$

But the second step in the above can also be written as,

$$\begin{aligned}
 &(GVA_1 + GVA_N) \times \frac{\sum_{i=1 \dots N} T(nul)_i}{(T(nul)_1 + T(nul)_N)} \\
 &= \frac{(GVA_1 + GVA_N)}{(T(nul)_1 + T(nul)_N)} \times \sum_{i=1 \dots N} T(nul)_i \\
 &= \frac{(GVA_1 + GVA_N)}{(T(nul)_1 + T(nul)_N)} \times \sum_{i \neq 1, N} T(nul)_i + (GVA_1 + GVA_N) \\
 &= \frac{(GVA_1 + GVA_N)}{(T(dat)_1 + T(dat)_N)} \times \frac{(T(dat)_1 + T(dat)_N)}{(T(nul)_1 + T(nul)_N)} \times \sum_{i \neq 1, N} T(nul)_i + (GVA_1 + GVA_N) \\
 &= \text{Weighted average margin} \\
 &\quad \times \text{Total turnover for non sampled firms} \\
 &\quad \times \text{Adjustment term to reflect change in average sample turnover} \\
 &\quad + \text{Total sampled GVA}
 \end{aligned}$$

This is equivalent to the way we are imputing GVA for non-sampled firms in our population dataset. In theory, our imputation methodology therefore applies the same weights to each non-sampled firm as the ONS’s grossing weights.

How we define each ‘bucket’ is also important here. Because the ABS survey is a stratified sample, for some combinations of four-digit sector, size and age group the ABS survey may only provide GVA data for a handful of firms. In these cases, the average GVA-turnover relationship for sampled firms may or not be a representative of non-sampled ones – there is simply no way of detecting whether this is the case. The consequence though could be undue volatility in the productivity estimates for non-sampled firms, which would introduce noise into our productivity decomposition.

To ensure that our firm-level productivity estimates are reasonable approximations of true productivity levels, for each firm we calculate four estimates of its GVA-turnover ratio, based on four definitions of the ‘bucket’ it belongs to. Each ‘bucket’ definition is progressively less granular on the SIC industry code dimension: the first ‘bucket’ is defined by its four-digit SIC code, size and age group, the second by its three-digit SIC code, size and age group, and so on. Our final estimate of GVA for each firm is an un-weighted average of the four GVA estimates generated by our four definitions of its ‘bucket’.

Tables D.1 and D.2 below shows how the sample and population datasets compare, both as aggregates and separated into age groups.

Table D.1: Comparison of sample and population by demographics (average between 2002 and 2011)								
	Sample				Population			
	Full	Births	Survivors	Death	Full	Births	Survivors	Death
Number of firms (000s)	452	29	403	18	17,182	1,853	13,450	1,556
Productivity (£000s /head/year):								
- Mean	4965	380	510	352	351	343	362	256
- Std. Deviation	1,243	631	1,295	678	819	433	894	442

Table D.2: Comparison of sample and population by age group (average per year between 2002 and 2011)										
	Sample					Population				
	Full	Age group				Full	Age group			
		1	2	3	4		1	2	3	4
Number of firms (000s)	452	68	89	263	32	17,182	3,924	4,445	6,465	2,348
Productivity (£000s /head/year):										
- Mean	496	468	534	502	363	351	380	374	345	252
- Std. Deviation	1,243	1,220	1,484	1,177	949	819	778	918	842	523
Age group 1: Firms in the first 2 full years of their lives										
Age group 2: Firms with age between 3 and 7										
Age group 3: Firms with age of 7 or above										
Age group 4: Firms within 2 years of dying										

Appendix E: Separating within-sector and across-sector reallocation effects

In section 6.4, the decomposition formula in BBH was modified to allow the reallocation term to be split into reallocation between firms in the same industry (defined by their four-digit SIC code) and reallocation between firms of different industries. The formula we applied can be found below.

$$\begin{aligned}
 \Delta \Pi_t = & \sum_{j \in Sectors} \sum_{i \in Surv} D_i^j \bar{\rho}_j \bar{\theta}_i \Delta \pi_{i,t} && \dots\dots \text{within firm} \\
 & + \sum_{j \in Sectors} \sum_{i \in Surv} D_i^j \bar{\rho}_j \Delta \theta_{i,t} (\bar{\pi}_t - \bar{\pi}_j) && \dots\dots \text{reallocation between surviving firms} \\
 & && \text{from the same sector} \\
 & + \sum_{j \in Sectors} \Delta \rho_{j,t} (\bar{\pi}_t - \bar{\Pi}^{Surv}) && \dots\dots \text{reallocation between different sector} \\
 & + \sum_{i \in Entry} \phi_{i,t} (\pi_{i,t} - \bar{\Pi}) && \dots\dots \text{reallocation to new firms} \\
 & - \sum_{i \in Exit} \phi_{i,t-1} (\pi_{i,t-1} - \bar{\Pi}) && \dots\dots \text{reallocation from exiting firms}
 \end{aligned}$$

where $\phi_{i,t}$ is firm i 's share of whole-economy employment at time t ,

$\pi_{i,t}$ is GVA per head in firm i at time t ,

Π_t is aggregate GVA per head at time t ,

$\rho_{j,t}$ is sector j 's share of survivors' employment at time t ,

$\theta_{i,t}$ is the share of firm i 's employment among surviving firms in its sector at time t ,

$\pi_{j,t}$ is GVA per head of survivors in sector j at time t ,

Π^{Surv} is aggregate GVA per head for all surviving firms at time t ,

D_i^j is a dummy variable that takes the value of 1 when firm i is located in sector j , and

a bar over a variable indicates the average of the variables across $t-1$ and t .

References

Arrowsmith, M., J. Franklin, Gregory, M. Griffiths, E. Wohlmann, G. Young, 2013, “SME forbearance and its implications for monetary and financial stability”, Bank of England Quarterly Bulletin, Vol.53, No.4, pages 296-303.

<http://www.bankofengland.co.uk/publications/Documents/quarterlybulletin/2013/qb130401.pdf>

Baily, M. N., E. J. Bartelsman and J. Haltiwanger, 2001, “Labor productivity: structural change and cyclical dynamics”, Review of Economics and Statistics, Vol. 83 (3), pages 420-433.

Bartelsman, E. J, J. Haltiwanger and S. Scarpetta, 2013, “Cross-country differences in productivity: the role of allocation and selection”, American Economic Review, American Economic Association, vol. 103(1), pages 305-34, February.

Basu, S., J. G. Fernald and M. Kimball, 2006, “Are technology improvements contractionary?”, American Economic Review, American Economic Association, vol. 96(5), pages 1418-1448, December.

Blundell, R., C. Crawford and W. Jin, 2013, “What can wages and employment tell us about the UK's productivity puzzle?”, IFS Working Papers W13/11, Institute for Fiscal Studies.

Broadbent, B., 2012, “Productivity and the allocation of resources”, Bank of England, Speech at Durham Business School.

<http://www.bankofengland.co.uk/publications/Documents/speeches/2012/speech599.pdf>

Caballero, R., 2007, “Specificity and Macroeconomics Restruction”, The MIT press.

Caballero, R. and E. Hammour, 2000, “Creative Destruction and Development: Institutions, Crises, and Restructuring”, NBER Working Papers 7849, National Bureau of Economic Research, Inc.

Caballero, R. and E. Hammour, 2005, “The Cost of Recessions Revisited: A Reverse-Liquidationist View”, Review of Economic Studies 72, pp. 313-341.

Crawford, C., W. Jin and H. Simpson, 2013, “Productivity, investment and profits during the Great Recession: evidence from UK firms and workers”, Fiscal Studies, vol. 34, no. 2, pp. 153-177.

Criscuolo, C., J. Haskel and R. Martin, 2003, “Building the evidence base for productivity policy using business data linking”, Centre for Research into Business Activity, Economic Trends 600, November.

<http://www.ceriba.org.uk/pub/CERIBA/BuildEvidenceBase/ETNov03Haskel.pdf>



Disney, R., J. Haskel and Y. Heden, 2003, “Restructuring and productivity growth in UK manufacturing”, *Economic Journal*, Vol. 113, pp.666-694.

Fay, J. A., and J. L. Medoff, 1985, “Labor and output over the business cycle: Some direct evidence”, *American Economic Review* 75 (no. 4), pages 638-655.

Fort T.C., J. Haltiwanger, R. S. Jarmin and J. Miranda, 2013, "How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size", *IMF Economic Review*, Palgrave Macmillan, vol. 61(3), pages 520-559, August.

Foster, L., C. Grim and J. Haltiwanger, 2013, “Reallocation in the Great Recession: cleansing or not?”, US Census Bureau Center for Economic Studies, Working Paper 13-42.

Foster, L., J. Haltiwanger and C. J. Krizan, 1998, “Aggregate productivity growth, lessons from microeconomic evidence”, NBER Working Paper 6803.

Goodridge, B., J. Haskel and G. Wallis, 2013, “Can intangible investment explain the UK productivity puzzle?”, *National Institute Economic Review*, No. 224, pp R48-R58.

Harris, R., 2002, “Foreign ownership and productivity in the United Kingdom - some issues when using the ARD establishment level data”, *Scottish Journal of Political Economy*, vol. 49, August 2002, pp 318–335.

Harris, R. and J. Moffat, 2013, “Total Factor Productivity Growth in Local Enterprise Partnership Regions in Britain, 1997-2008”, *Regional Studies*, Published on-line, DOI: 10.1080/00343404.2013.799770.

Hughes, A. and J. Saleheen, 2012, “UK labour productivity since the onset of the crisis — an international and historical perspective”, *Bank of England Quarterly Bulletin*, Vol. 52, No. 2, pages 138-146.

Levine, O. and M. Warusawitharana, 2014, "Finance and Productivity Growth: Firm-level Evidence", *Finance and Economics Discussion Series* 2014-17, Board of Governors of the Federal Reserve System.

Monetary Policy Committee, 2013, Minutes of the MPC Meeting held in January and February.

Office for National Statistics, 2012, “International comparison of productivity- Final Estimates for 2011, available at <http://www.ons.gov.uk/ons/rel/icp/international-comparisons-of-productivity/2011---final-estimates/stb-icp-final-estimates-2011.html>

Oi, W., 1962, “Labor as a Quasi-Fixed Factor”, *Journal of Political Economy*, December, 70(6), pp. 538-55.

Oulton, N. and M. Sebastia-Barriel, 2013, “Long and short-term effects of the financial crisis on labour productivity, capital and output”, Bank of England Working Papers No.470.

Riley, R., C. Rosazza Bondibene and G. Young, 2013, “ Financial Crisis and UK Company Performance”, mimeo

White, K., J. Reiter and A. Petrin, 2012, “Plant-level Productivity and Imputation of Missing Data in U.S. Census Manufacturing Data”, NBER Working Paper No. 17816

