



BANK OF ENGLAND

Speech

The UK's Productivity Problem: Hub No Spokes

Speech given by

Andrew G Haldane, Chief Economist, Bank of England

Academy of Social Sciences Annual Lecture, London

28 June 2018

The views expressed here are not necessarily those of the Bank of England or the Monetary Policy Committee. I would like to thank Sandra Batten, Shiv Chowla, Gabor Pinter, Patrick Schneider, Brad Speigner and Pawel Smietanka for their help in preparing the text. I would like to thank Tera Allas, Ben Broadbent, Caroline Butler, Jeremy Franklin, Colm Manning, Richard Sherrington, Silvana Tenreyro, Ryland Thomas and Jan Vlieghe for their comments and contributions.

This work contains statistical data from 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 datasets which may not exactly reproduce National Statistics aggregates.

Introduction

I am delighted to be giving this year's Annual Lecture of the Academy of Social Sciences and honoured to be a Fellow of the Academy. I thought I would use the occasion to discuss an issue of pressing public policy concern, an issue that affects the living standards of each and every one of us and has, as best we can tell, since the dawn of time: productivity.

It is a terrible word, as it leaves most people dazed and confused. Few are those who can define it and fewer still those who can measure it. Yet it has entered the popular lexicon and with good reason: the one thing we do know is that productivity is crucial to our pay and living standards over the longer run. Productivity is what pays for pay rises. And productivity is what puts the life into living standards.

History, distant and recent, makes this only too clear. It is no coincidence that productivity, pay and living standards have moved in lock-step over the past millennium – flat-lining for the first three-quarters before sky-rocketing in the final quarter - as each new industrial revolution has dawned (Chart 1). Pay and living standards have increased over fifteen-fold since 1750. Productivity explains why.

More recent experience has been more painful. Since 2008, productivity in the UK has essentially flat-lined. This is almost unprecedented in the modern era, a “lost decade” and counting. And it has been mirrored, as in the past, in a lost decade for pay too (Chart 2). Inflation-adjusted pay has stood still since the crisis, something also almost unprecedented in the modern era.

This prolonged pause has become known as the productivity problem.¹ In the UK, the problem is a big one by any historical standard. And it is cold comfort that the UK shares this problem with much of the Western World and a decent chunk of the emerging market world. Cold comfort because the UK's productivity slowdown appears to have been larger than in almost any other country.

Given its seriousness, it is no surprise there has been no shortage of words written, and proposals made, on how to understand and solve the productivity problem.² If productivity was measured by words written and spoken on the subject, it would have long-since been solved. The UK government has issued its own green and white papers, setting out an industrial strategy.³ These policy proposals hold, I believe, real promise.

For today, I thought I would begin by providing some diagnostic evidence on the UK's productivity problem, focusing on the long and lengthening tail of lower-productivity firms.⁴ This is not the only lens through which to view the UK's productivity problem. Looking at the problem by sector, by region or by city all have merit

¹ Recent references to the 'UK productivity problem' include Macpherson (2018) and Wren-Lewis (2017).

² For example, Aghion *et al* (2017), Jones (2016), Jacobs *et al* (2017).

³ HM Government (2017) and Department for Business, Energy & Industrial Strategy (2017).

⁴ Haldane (2017), Berlingieri *et al* (2017), Andrews, Criscuolo and Gal (2015), Andrews, Criscuolo and Gal (2016).

too.⁵ But I find the distributional lens a useful one for identifying the source of the UK's "productivity gaps" – its gap with competitor countries and its gaps with its own past performance.

It also leads naturally to a discussion of why those gaps have emerged between the best and the rest of UK companies and what might be done to close them. Can ways be found to accelerate the rate of technological diffusion to the long tail? I will argue an improved *diffusion infrastructure*, using as a base existing institutions and the UK's innovation strengths, could help to close those productivity gaps.

The "Hub-No-Spokes" Model

My diagnosis of the UK's productivity problem starts with three simple stylised facts about productivity gaps. The first two gaps are well-known, the third much less so. Fact one is that UK productivity has flat-lined for a decade. This means UK productivity is running almost 20% below its level had it continued along its pre-crisis trend (Chart 3). This is a big productivity gap, even if you believed the pre-crisis productivity trend may have been somewhat over-inflated.⁶ It is a gap it is unclear will ever be closed.

Fact two is that there is a second gap – perhaps as much as a third – between levels of productivity in the UK and in our main competitors, the US, Germany and France (Chart 4). This is an even-larger gap, even if you believed productivity was differently (mis-)measured across countries. It is the source of the oft-quoted quip that the average French worker achieves by Thursday lunchtime what the average British worker achieves only by close of business on a Friday.

Fact three is that the productivity gap between the top- and bottom-performing companies is materially larger in the UK than in France, Germany or the US. In the services sector, the gap between the top- and bottom-performing 10% of companies is 80% larger in the UK than in our international competitors (Chart 5).⁷ This productivity gap has also widened by far more since the crisis – around 2-3 times more – in the UK than elsewhere (Chart 6).⁸

This third gap is interesting, in part, because it is striking and less well-known. But, more fundamentally, it is interesting because it can help in explaining the first two gaps. A long and lengthening tail of stationary companies explains why the UK has a one third productivity gap with its international competitors. And that same tail is also an important reason why the UK has a one fifth productivity gap relative to the past.

⁵ For example, Tenreyro (2018), Swinney (2018), Remes *et al* (2018). An alternative strand of the literature considers the role of market power, for example De Loecker and Eeckhout (2017). There is also a strand that debates mismeasurement of GDP and productivity, for example, Mokyr (2014).

⁶ For example, Brennan, Haldane and Madouros (2010).

⁷ The gap in manufacturing is around 50%.

⁸ UK estimates of firm-level productivity dispersion are based on Bank of England analysis of ONS data, whereas non-UK estimates of productivity dispersion are provided by the OECD. Differences in underlying data used may therefore play a role in Charts 5 and 6.

To mix metaphors for a moment, the UK is in many respects a tale (tail) of two companies: a small set in the upper tail gazelling along the productivity high road and a much larger set in the lower tail snailing along the low road. Capitalism always of course throws up winners and losers, hares and tortoises, gazelles and snails, upper and lower tails. In the UK, however, these differences are far-larger, and have increased by more, than elsewhere. And therein, I would argue, lie the roots of the UK's productivity gaps.

Let me start with the good news, focussing on companies on the high road. The first and most important point to make is that the UK does not lack for innovative, high-productivity companies. If you tour the UK, as I do, you find these "frontier" companies in every region and every sector. Their productivity is world-leading, their technology world-beating, their managers and workers world-renowned. They are inspirational.

Micro-analysis of UK companies confirms this anecdotal impression. Chart 7 plots the distribution of UK relative to French and German companies.⁹ There are more UK companies in the upper tail than in Germany, with its much-vaunted industrial reputation, and France. The top 10% of UK companies have levels of productivity at least 100% above the median. The top-decile of German and French companies have, by comparison, levels of productivity 80% and 60%, respectively, above the median.

If we look more closely at companies in the top 1% or 0.1% of the productivity distribution, the scale of this success story becomes clearer still (Table 1). On average over the ten years to 2014, 1%-er companies experienced annual productivity growth of 8%. 0.1%-er companies have seen annual productivity growth of 12%. There is no evidence of a lost decade among these high-fliers. For them, the Fourth Industrial Revolution has already arrived.

It is tempting to think these frontier companies are a tightly-knit bunch, geographically and sectorally. By and large, you should resist that temptation. Table 1 shows productivity growth over the past decade among the top 1% and 0.1% of companies, broken down by region and by sector.¹⁰ There are some important differences in productivity performance. Even among the 0.1%-ers, there are some wide regional variations in performance, with the East Midlands at the top (30% per year) and Wales at the bottom (0%).

Differences across sectors are larger still. The upper tail of sectors such as transportation and storage and information has seen annual rates of productivity growth of 30 or even 40%. Among virtually all sectors, and across virtually all regions, however, upper tail companies have seen positive and healthy rates of productivity growth. The Fourth Industrial Revolution has reached almost all sectors and all regions to some degree, even as the productivity problem in the UK as a whole has perpetuated.

⁹ These numbers control for differences in the sectoral composition and sizes of firms between the countries in question. Further detail of the method used here is set out below Chart 7 in the annex.

¹⁰ Data refer to non-financial business sector.

So how have these upper tail companies been able to buck the UK trend? The short answer appears to be because the UK is, on many dimensions, a global innovation hub. At a headline level, various indices of countries' innovation status have been constructed. The UK has appeared in the top 5 of the Global Innovation Index for each of the past five years, ahead of France and Germany.¹¹ It also ranks ahead of them on the European Innovation Scorecard; it describes the UK as an “innovation leader”.

Digging a little deeper, the UK scores particularly highly on research and development and start-ups. For example, it tops the charts among rankings of the world's top universities, with 5 of the world's top 25 universities according to *The Times* rankings for 2018, including the top two slots. And it also ranks first globally in categories such as citable scientific publications, e-commerce and ICT. It ranks fourth globally for its creative outputs. The UK is the largest magnet for tech talent in Europe.

This creativity and entrepreneurship is reflected in business start-ups. Over 1,100 new businesses start up in Britain each *day* – roughly, one every 75 seconds.¹² And when started, a disproportionately large number succeed, with the UK home to half of the top 10 fastest-growing companies in Europe.¹³ Financing them is a venture capital industry which, in 2017, invested more in the UK than in Germany and France combined.

The upper tail of UK companies, then, is large and thriving. It more than holds its own relative to international competitors. Although there is ample scope for improvement, the upper tail is fairly well-spread, sectorally and geographically. Its success is built on world-leading education, innovation and research. The UK is, genuinely, an innovation hub, with companies to match.

Let us now reduce altitude and take a closer look at the lower tail. As Chart 7 also illustrates, there is a materially larger lower tail of UK companies, productivity-wise, than there are in Germany and France.¹⁴ The bottom 25% of UK companies have levels of productivity around 80% or more below the UK median. Their Germany and French counterparts have productivity around 60% or more below the median.

If we look more closely at this lower tail, what is striking is both its scale and longevity. Table 1 shows growth in average productivity among the bottom 99% of the productivity distribution. Since 2004 – that is, before the crisis – average productivity growth among the 99%-ers has averaged only 1%. This suggests that, among a significant swathe of UK companies, levels of productivity must have flat-lined or fallen.

This lower tail of companies shares one important feature with the upper tail: neither is confined to particular regions or sectors. Near-zero rates of productivity growth among the 99%-ers are a feature of all regions

¹¹ Cornell University, INSEAD and WIPO (2017).

¹² Department for Business, Energy & Industrial Strategy (2017).

¹³ Department for Business, Energy & Industrial Strategy (2017).

¹⁴ These numbers control for differences in the sectoral composition and sizes of firms between the countries in question. Further detail of the method used here is set out below Chart 7 in the annex.

and almost all sectors (Table 1). Even the best-performing region (London) and best-performing sector (professional, scientific and technical) in the lower tail has mustered productivity growth of only 2% and 4%.

We can look more precisely at the fraction of firms who have seen effectively no rise in productivity since before the crisis.¹⁵ For the UK as a whole, that amounts to around half of them. For some regions and sectors, the fraction of firms that have stood still in the 21st century is closer to two-thirds. This tail of serially-stagnating companies is long indeed.

It also goes quite some way towards accounting for the UK's two productivity gaps. It is, arithmetically at least, the reason why the UK has a productivity gap with France and Germany. The positive effect of the UK's larger upper tail of high-performing companies is being comfortably offset by the negative effects of its much longer, lower tail. The UK's international productivity gap is, to a large degree, a long tail problem.

This long tail long pre-dates the crisis and thus cannot by itself explain all of the UK's poor productivity performance since the crisis. But as my colleague Patrick Schneider has shown, a striking feature of the past few years has been how previously high-performing companies have seen their productivity laid-low.¹⁶ These are companies which, previously, were nestled just inside the technological frontier.

The crisis has caused these companies to stand still, productivity-wise, while frontier companies have forged ahead. Companies previously gliding in the slipstream are now floundering in the wake of frontier firms. That has widened the distance between the best and the rest; it has lengthened the lower tail. And that lengthening helps explain the UK's productivity slowdown over the past decade.

There is an active debate underway, in academic and policy circles, about the forces shaping future productivity and growth. Often, this debate centres on whether the forces of secular innovation, brought about by the dawning of the Fourth Industrial Revolution, will win out over the forces of secular stagnation, arising from adverse trends in demography, inequality and the like.¹⁷

The UK is, in many respects, a prime example of that secular struggle in practice. The forces of secular innovation and secular stagnation are both operating powerfully, with the two productivity poles acting like magnets pulling in opposite directions. This is widening the distribution of productivity across firms and, as a by-product, contributing to a widening of the distribution of pay across workers. The tail of two companies has, as its mirror-image, a tail of two workers.

Another perspective on these productivity dynamics is provided by viewing them through the lens of technological diffusion. In the fullness of time, innovation should be expected to diffuse through the

¹⁵ Defined as the share of firms with annualised productivity growth of less than ½% between 2004 and 2014.

¹⁶ Schneider (2018).

¹⁷ For example, Gordon (2012) and Brynjolfsson and McAfee (2014).

economy, lifting all boats. That has been the lesson of every industrial revolution. Yet in the UK, this technological trickle-down, from frontier to tail, appears to have dried up. A lengthening flotilla of boats has remained in dry dock. The diffusion engine appears, for them, to have seized up.¹⁸

Empirical evidence, using international innovation indices, supports this diagnosis. Despite ranking consistently in the top-5 for innovation, the UK ranks only 38th globally for knowledge diffusion.¹⁹ This is far lower even than England's ranking in the World Cup.²⁰ Moreover, that diffusion ranking has fallen very sharply over recent years, again in lockstep with England's world football ranking.²¹

It is not just that the UK has a lowly ranking on R&D expenditure by international comparison – at 1.7% of GDP, more than 1 percentage point below our main competitors.²² It is that three-quarters of the UK's private R&D spending is concentrated in only 400 companies, less than 0.01% of the UK's business population.²³ Among smaller firms, the UK ranks low down the European league table on in-house innovation and the introduction of new products and processes.²⁴

Typically, we think of “Research and Development” (R&D) as a rhyming couplet. In the UK's case, the R and the D do not seem to rhyme. The UK does R well, as a world-leading innovation hub. But it does D poorly, where the D refers not just to development but the diffusion and dissemination of innovation to the long, lengthening, languishing lower tail. When it comes to innovation, the UK is a hub without spokes.

Understanding the Barriers to Diffusion

What explains why the UK is poor at disseminating and diffusing the innovation it is rich at generating? Answering that question is important for identifying potential barriers to improved diffusion and, hence, for devising policies for breaking down these barriers and supporting productivity growth.

A useful framework for thinking about the determinants of technological diffusion has been provided by economists Diego Comin and Marti Mestieri.²⁵ Drawing on evidence from a wide range of countries and technologies over many centuries, they highlight four structural factors which are crucial for technological trickle-down within and across economies. I shall consider the UK evidence on each in turn.

¹⁸ To a lesser extent, there is some evidence of the same being true internationally (Berlingieri *et al* (2017)).

¹⁹ Cornell University, INSEAD and WIPO (2017).

²⁰ Which currently stands at 12th, but just you wait...

²¹ Between 2013 and 2017, the UK fell from 18th to 38th in the global ranking for knowledge diffusion. Since 2012, England have fallen from 6th to 12th in the FIFA world rankings.

²² Department for Business, Energy & Industrial Strategy (2017).

²³ Department for Business, Energy & Industrial Strategy (2017). According to Rhodes (2017), the UK is estimated to have had 5.7 million private sector businesses.

²⁴ European Commission (2018).

²⁵ Comin and Mestieri (2013). We use a slight variant of Comin and Mestieri's organising framework. For example, we do not consider the influence of demand expectations on innovation and diffusion - a shorter-term cyclical factor affecting firms' investment plans and capital-deepening. Teneyro (2018) discusses the contribution of reduced capital-deepening to the UK's post-crisis productivity performance.

(a) *External Openness*

Over recent decades, globalisation has resulted in secular rises in cross-border flows of goods and services, capital, people and information. One of the key potential benefits of this sweeping globalisation and openness is that it can serve as a vehicle for the transmission and diffusion of new technologies across borders.²⁶ Globalisation would tend to accelerate catch-up between countries operating at the technological frontier and those operating inside.²⁷

The evidence appears to support the first leg of this hypothesis: increased openness, across its various dimensions, does appear to be associated with higher levels of productivity as innovation diffuses. To illustrate, Charts 8-11 show some indices of external openness (for goods and services, capital, people and a summary of all three) plotted against levels of productivity in a selection of 32 countries.²⁸

The relationship between them is positively sloped and statistically significant: countries with greater external openness - whether trade, capital or people – tend to have higher levels of productivity. It would be brave to claim causality. But it is consistent with openness of factor flows being an important cross-border conduit for ideas and innovation, as indeed other studies have found.²⁹

Adding weight to that, what is true *across* countries appears also to be true *within* countries. Charts 12 and 13 plot the productivity distribution of UK companies sliced in two different ways. Chart 12 distinguishes companies according to whether or not they export. Exported-oriented companies have average levels of productivity one-third higher than their non-exporting counterparts.

Chart 13 looks at a different dimension of external openness, namely whether a company is foreign-owned. The performance of foreign-owned firms captures some of the wider benefits of cross-border flows (foreign capital, ideas and people). The productivity differences with domestically-owned firms are even more striking. Productivity is around twice as high in foreign-owned firms.³⁰

These stylised facts suggest external openness is an important driver of productivity. Whether it can explain the UK's productivity gaps is, however, more questionable. The UK in fact scores highly on all the external openness metrics. Given that degree of openness, its productivity under-performs other countries; it sits beneath the regression line of best-fit (Charts 8-11). Openness cannot, on the face of it, easily explain the UK's productivity gap with other countries.

²⁶ For example, Feyrer (2009a) and Feyrer (2009b).

²⁷ Baddeley (2006).

²⁸ The indices of openness come courtesy of the McKinsey's Global Institute as shown in Manyika *et al* (2016).

²⁹ For example, Frankel and Romer (1999), Feyrer (2009a) and Feyrer (2009b).

³⁰ Batten and Jacobs (2017).

Nor can it explain the UK's productivity slump since the crisis. Over recent years, most measures of the UK's external openness have been increasing, not decreasing. Cross-border flows of people, direct investment and trade into the UK have risen to their highest levels in a generation in the past few years, while flows of information have risen to their highest-ever levels, at the same time productivity has flat-lined.

A similar openness puzzle is operating at a global level. Cross-border flows of goods and services, money, people and especially information have all risen to their highest levels, in some cases, ever. Yet at the same time, levels of productivity across countries have tended to *diverge*. So too, as a result, have levels of GDP per head. Technological trickle-down and GDP catch-up has gone into reverse gear at the same time external openness has switched up a gear. How do we explain this?

(b) *Technology Transfer*

To begin doing so, it is useful to recognise that two distinct forces are at work when technology is transferring across countries or companies.³¹ First, there is *technological adoption* – the time it takes a new technology to first reach a country or company. Second, there is *technological penetration* – the extent to which technology then reshapes processes and products in a company or country.³²

Economists often refer to the first of these as the *extensive margin*, the second as the *intensive margin*. In principle, we might expect these two aspects of technology transfer to move in lockstep. The two dimensions to diffusion should be complementary. In practice, international evidence seems to suggest a puzzling divergence in the two dimensions of technology transfer across countries.

Average adoption lags for new technologies have tended to fall over time, consistent with rising cross-border flows of ideas, capital, goods and services and people. For example, the adoption lag for steam ships (18th century) was over a century, for the telephone (19th century) was around 50 years, while for the kidney transplant (20th century) was just 13 years.³³ Each century has seen a halving of adoption half-lives. 21st century technologies, like 3-D printing and AI, appear to have had a shorter-still adoption lag.

Over the same period, however, the intensity of use of these new technologies appears to have diverged between countries. For example, rates of technological penetration of steam ships, telephones and kidney transplants have shown a greater divergence between frontier and non-frontier countries with each passing century. It will be interesting, in time, to see if the same is true of 3D-printing and AI.

³¹ There is a particular literature which discusses the technology frontier where firms in regions/countries that are skilled-labour abundant tend to choose technologies that are best suited to skilled workers; firms in poorer regions/countries, which tend to be unskilled-labour abundant, choose technologies more appropriate to unskilled workers. Noteworthy papers include Acemoglu (2002), Acemoglu, Gancia and Zilibotti (2015) and Caselli and Coleman II (2006).

³² Comin and Mestieri (2013). Filippetti and Archibugi (2011) also discuss innovation across countries.

³³ Comin and Mestieri (2013).

What has been the net effect of these two opposing forces? In practice, it seems that cross-country divergences in rates of penetration have more than counter-balanced convergence in rates of adoption. And that, at least in an arithmetic sense, explains why we have seen a puzzling divergence in levels of productivity, and living standards, across countries since the 1980s.³⁴

Globalisation has meant that ideas and innovations are indeed spreading more rapidly and more widely than perhaps ever previously. But for later-adopters of these technologies, they are spread thinly. And despite thickening over time, penetration of these technologies among later-adopters never appears to catch up. Nor, then, do levels of productivity and living standards.

If that is true *across* countries, might it also be true *within* countries? Do the differing dimensions of technology transfer help explain the slowing rates of diffusion from frontier companies to the long tail? There is no ideal economy-wide measure of diffusion and patterns differ technology by technology. Nonetheless, it is possible to build up a reasonable picture by looking at transfer patterns among some specific technologies.

Looking first at adoption rates, on the face of it the UK fares reasonably well looking across a range of technologies. Firms with internet access and a website are high in the UK by international comparison. Chart 14 looks at the number of UK firms adopting five basic ICT technologies (computer, internet, website, e-purchases and e-sales) on three dates (2002, 2007, 2015). There has been a clear and significant rightwards shift over time in this distribution, consistent with increasing adoption.

Nonetheless, by 2015 there were still only 13% of companies who had adopted all five of these basic technologies. And there was still a tail of 9-10% of companies who had adopted only 2 or fewer of them. This tail appears to be largely made up of smaller firms. In innovation surveys, UK SMEs fare poorly by European comparison on introducing new product and process innovations, ranking 20th out of 36.³⁵

This picture is less encouraging still when it comes to the adoption of more advanced technologies. Chart 15 looks at adoption rates across UK companies of six such technologies (mobile access to email, documents and software, websites with online ordering, fast broadband access and electronic data interchange sales) on two dates (2011 and 2015).

Progress towards adoption of these technologies appears to have been slow between 2011 and 2015, with only a modest rightward shift. Even by 2015, fewer than 10% of UK companies had four or more of these advanced technologies. More than half of all UK companies had adopted one or less of them and a quarter had yet to adopt any. This suggests a patchy picture of technological adoption among UK companies, with a long tail of slow-adopters of basic and especially advanced technologies.

³⁴ Haldane (2017), Comin and Mestieri (2013).

³⁵ European Commission (2018).

If we turn to penetration rates, the picture is in some respects similar. There appears to be a significant but small upper tail of companies with high utilisation of technologies and a large lower tail whose penetration rates are low and patchy. Chart 16 illustrates this for one technology - computing. It shows the percentage of employees in a company using a computer on three dates (2002, 2008, 2015). The good news is that it shows a clear trend towards increased employee computer use over time.

The distribution of computer use is, however, clearly twin-peaked. There is an upper tail of companies where more than 90% of staff use computers. By 2015, this had increased to a little under half of all companies. But there is also a large, if diminishing, tail of companies where employee use of computers is much more modest. For example, even by 2015 around 40% of companies had fewer than half of their employees using computers. And there were still 20% of companies where fewer than one in five employees used computers.

This picture is patchier still for more advanced technologies. Chart 17 looks at the distribution of the proportion of employees in UK companies with access to a portable device with internet access on two dates (2011 and 2015). It, too, is twin-peaked. But the upper peak is now a very modest one, containing only around 10% of firms. And the lower peak, by contrast, is Himalayan. The fraction of companies where 50% or less of employees use portable devices is around 80%.

Patterns of technological adoption and penetration among UK companies, then, tell a broadly consistent story. There are clearly identified upper-tail companies, who are fast adopters of new technologies which then appear to spread quickly and widely across the firm. Adoption of innovation is rapid and diffusion widespread. This helps explain why these companies are on the productivity high road.

But there are many more companies where, by contrast, technological adoption has been slow or non-existent. Even when adopted, technologies have tended not to have been quickly or extensively employed. For this lower tail, neither innovation nor diffusion has been rapid or widespread. This, in turn, probably helps explain why these companies' productivity performance has been stuck on the low road.

(c) *Human Capital*

One of the key mechanisms through which ideas and innovations are diffused across countries and companies is, unsurprisingly, through people. As workers transition between companies, their expertise and experience is transferred within them. Empirical evidence from a range of countries confirms the importance of human capital as a transmission and diffusion channel.³⁶

Can these human capital channels help explain the UK's productivity gaps? One aspect studied in some depth is management skills. A survey by Nick Bloom and John van Reenen found that UK companies had

³⁶ For example, Davis and Haltiwanger (2014) and Bjelland *et al* (2011).

management skills around half a standard deviation lower than competitor countries, such as the US. These management skills-scores have a statistically-significant link with levels of company productivity.³⁷

The UK also exhibits a greater degree of dispersion in management skills than competitor countries. Chart 18 plots the distribution of management skills from the survey in the UK, Germany and the United States. The UK distribution has a notably larger lower tail. The UK has around twice as many companies with a management score below 2 than in Germany and the US, consistent with the UK's longer tail being sourced, at least in part, in weaker management skills.

Recent research by Bank colleagues, Saleem Bahaj, Angus Foulis, and Gabor Pinter, suggests this tail may also be explained by skills in corporate board rooms. Using company account data, they are able to assess the degree of "connectivity" of board members to other companies. This connectivity proxies a company's ability, through its board, to draw on the experience of other companies when learning about and adopting new practices, products and processes.

Chart 19 and Table 2 plots this measure of board connectivity across UK companies. It exhibits the same "power-law" distribution as other measures of connectivity, with a small set of companies whose directors are hyper-connected and a large set of firms whose directors are loosely or unconnected.³⁸ As with Facebook friends and Twitter followers, there is a very small "in-crowd" of well-connected companies and a large disconnected "out-crowd".

We can go one step further and ask whether this long-tailed distribution of board connectivity has any systematic relationship with the long tail of companies' productivity performance. Controlling for other firm-specific factors, such as their size and sector, it appears that it does. Indeed, this relationship is both statistically significant and economically meaningful (Table 3).

As a thought-experiment, imagine comparing two otherwise-identical companies except for their degree of board connectivity. One company has a board whose members are in the upper decile of the connectivity distribution; it is one of the "in-crowd". Another has a board drawn from the lowest connectivity decile; it is one of the majority in the "out-crowd". Average annual productivity growth for the in-crowd company is around 1 percentage point higher than for the out-crowd company.

I am indebted to the doyen of monetary and financial economics, Professor Charles Goodhart, for many things. But among them is the following, the only corporate governance joke I know. Question: what is the difference between a shopping trolley and a non-executive director of a company? Answer: a shopping trolley has a mind of its own. The evidence here rejects the null hypothesis in Charles's joke.

³⁷ Bloom and van Reenen (2007), Haldane (2017).

³⁸ Watts (2003), Haldane (2009), Ferguson (2018).

It is unclear whether, for Facebook fanatics and Twitter addicts, higher numbers of friends and followers translates into higher levels of productivity and well-being. My casual empiricism suggests, sometimes, it might not. But for companies, the empirical evidence suggests connections do matter. For companies' productivity, connectivity counts. Indeed, it may help account for the UK's long tail.

Another window on the UK's productivity gaps is provided, not by looking at the *stock* of human capital (such as in management or boards), but at the *flow* of these skills between companies as workers move jobs. So the theory goes, the larger those worker flows – so-called labour market “churn” – the greater ought to be the transfer of knowledge and skills and the faster the pace of technological diffusion.

Chart 20 plots a measure of labour market churn in the UK, with the US as a point of comparison. Two points stand out. Labour market churn rates are systematically higher in the US than in the UK, by a factor of around 2.³⁹ This is likely to have contributed to slower rates of technological diffusion in the UK than US. It is also consistent with the UK's significant productivity gap with the US.

A second observation is that turnover in the UK labour market has been running well below its historic levels for much of the past decade. This probably reflects the impact of heightened job insecurity on people's willingness to move firm in the aftermath of the crisis. One of the adverse side-effects has been a slowing in the flow of knowledge and ideas between firms, potentially gumming up the technology-transfer process.

We can look at these churn rates on a more disaggregated basis by linking data on worker movements and company performance.⁴⁰ Chart 21 plots the evolution of churn rates for firms above and below the productivity median. Pre-crisis churn rates were systematically (around 20%) higher in high-productivity firms than in low-productivity ones. Ideas and skills recycled more quickly in the upper than in the lower tiers of the productivity distribution.

This may have contributed to the long tail problem. Other things equal, one might wish to see higher rates of labour market churn among poor-performing companies as a means of injecting new ideas, experience and skills, speeding-up the process of technological diffusion to them. In practice, the opposite appears to have been the case. The lower tail of UK companies, pre-crisis, had a stickier workforce.

Post-crisis, things have changed. Churn rates among high and low productivity firms have converged, with churn rates among higher productivity companies falling sharply. This matches the productivity patterns seen post-crisis, with the slowdown focussed on previously high-performing companies. Slower churn among these companies, and the accompanying slower diffusion, may have contributed to their slowing productivity and them becoming part of that lower tail.

³⁹ The US job-to-job flow data are more comprehensive than those in the UK at picking up job moves within a given quarter. This methodological difference means that estimated US churn rates will be stronger than those in the UK, but this can only explain part of the gap between the UK and US job-to-job transition rates.

⁴⁰ We combine the ASHE and BSD data for this analysis.

A different way of slicing these granular data is to look at where workers are moving to. Again in an ideal world, we might wish to see high rates of worker transition to different points in the productivity distribution, to different sectors and different regions. Indeed, in a perfect world we might wish to see workers moving from higher-productivity to lower-productivity firms to enable good practices and processes to be effectively diffused from the upper to the lower tail.

Table 4 shows a matrix of transition probabilities for workers between four quantiles of the probability distribution for companies.⁴¹ All of these probabilities are conditional on a worker moving in the first place and so sum horizontally to one. Two points stand out. First, there is a high degree of stickiness in job moves. Workers tend to move between jobs in the same productivity quartile. In the upper and lower quartiles, around 40% of jobs moves are within-quartile.

The second point is that there is rather limited evidence of technological diffusion being supported by jobs flows from firms in the frontier to the long tail. The probability of a worker for a company in the upper quartile moving to one in the lowest quartile is around 20%. The relative dearth of such moves may help explain the lack of technological trickle-down to the long tail.

Taken together, the evidence here points to a two-tailed labour market. The upper tail in-crowd attracts the most skilled managers and the best-connected boards, amplifying their strengths. And in-crowd workers swim mostly in the pool of in-crowd companies. By contrast, out-crowd companies in the lower tail have less-skilled managers, less connected boards and workers who swim largely in less productive company pools. This bifurcated labour market seems unlikely to support technological diffusion to the long tail.

(d) *Institutional Infrastructure*

There is a well-established school of thought linking improvements in the quantity and quality of institutions to rises in living standards. Countries with high-quality institutions tend, on average, to grow more rapidly and to have higher levels of productivity and living standards. Those without tend to be subject to larger and more frequent breakdowns in economic, financial and social order, which holds back significantly their growth, productivity and living standards.⁴² Institutions matter.

One of the channels through which institutions support growth is by nurturing investment and innovation. For example, upholding the rule of law and enforcing property rights is typically found to be crucial in supporting business investment and innovation.⁴³ Having those rights well-established reduces concerns among firms

⁴¹ As the data show suspiciously large rates of job churn between the bottom and top deciles of the firm-level productivity distribution, we strip those deciles out from our analysis here, and only look at the 2nd to 8th deciles. Our description of 'quartiles' here is therefore not strictly the quartiles of the complete distribution.

⁴² For example, North (1990) and Haldane (2017).

⁴³ Coase (1960).

that their assets might otherwise be expropriated, either by the state or another company. Security of property rights encourages a mass flourishing of investment and entrepreneurship.⁴⁴

Historically, companies' assets have tended to be physical and tangible, such as plant and machinery. But businesses and their assets are changing fast. Increasingly, businesses' assets are taking on an intangible form, such as intellectual property, patents and goodwill.⁴⁵ On some estimates, intangible assets are already larger than tangibles among UK and US companies. Protecting these intangibles from expropriation is every bit as important as for tangible assets.

Fortunately, in international league tables the UK tends to fare favourably on those fronts, with a well-established rule of law and property rights, supporting investment and innovation.⁴⁶ There are other aspects of UK's institutional infrastructure which fare less well, however, in ways which might help to explain its productivity gaps. Let me focus on two: its *diffusion* infrastructure and *financial* infrastructure. A simple comparison of UK and German institutional infrastructures is revealing here.

Germany has a long-established infrastructure to support innovation and its diffusion to German companies. One important component of this is the German *Fraunhofer* institutes. These were first established in 1948 to rebuild the German corporate sector. They have now grown to total 72 in number, covering all sectors and regions across Germany. Collectively, they employ almost 25,000 people and help around 6,000-8,000 companies each year, large and small.⁴⁷ They have an annual budget of a little under 0.1% of German GDP.

A second piece of infrastructure supporting German business is the so-called *Steinbeis* system. This was created in 1971. It is a network of technical professionals whose skills, experience and know-how can be drawn on by companies across Germany, operating through around 1,000 Steinbeis Enterprises. It currently has around 6,000 professionals in the network. Over 2,000 businesses have been founded on the back of the Steinbeis platform.

A third layer of the institutional cake in Germany is the financing of companies. This itself has three layers: the *Sparkassen*, operating locally and largely financing local businesses; the *Landesbanken*, operating regionally and typically serving somewhat larger companies; and *KfW*, Germany's national development bank, typically serving large companies. Often, representatives from these organisations will have a role in the governance of German businesses.

Although this description does not do justice to all of the institutional infrastructure supporting German companies, these three elements have played an important role in fostering innovation and its diffusion across German businesses: by promulgating ideas and technologies (through the *Fraunhofer*); by recycling

⁴⁴ Acemoglu and Robinson (2012).

⁴⁵ For example, Haskel and Westlake (2017).

⁴⁶ Though, as Haskel and Westlake (2017) point out, there is more that could be done to support intangible assets in businesses.

⁴⁷ See <https://www.fraunhofer.de/en/about-fraunhofer/profile-structure/facts-and-figures.html> and Comin, Trumbull and Yang (2011).

expertise and know-how (through the Steinbeis); and by supporting company financing and governance (through the Sparkassen, Landesbanken, KfW).

If we turn to the UK, it appears to have fewer of these institutional spokes supporting technological trickle-down to companies. In 2011 following a review by Hermann Hauser, the UK set up its own equivalent of the Fraunhofer – so-called Catapult Centres. The aim was to create a new innovation infrastructure across the UK. This has been a real success. Having visited a number of them, I can attest to the great work the Catapults are doing to promote innovation, acting in partnership with some of the leading UK companies.

The scale and scope of activities by the UK Catapult Centres, however, falls short of their Fraunhofer counterparts (Table 5). The Catapults are far fewer in number, at around 10. And they are significantly smaller in scale, with around 1,500 employees and an annual budget of just 0.01% of GDP. Their numbers of company collaborations, at around 600 each year, is more limited, especially to smaller companies.

The Catapult Centres have created a new innovation engine for the UK. This is great news. The less good news is that this engine has fewer cylinders than its Germany counterpart. Indeed, it serves more as an innovation than a diffusion engine, leaving largely untouched the long tail of UK companies. Nor does the UK have the equivalent of the Steinbeis system for recycling technical expertise and know-how between companies, which could help de-segment the skills-stratified UK labour market.

The financial infrastructure supporting companies is also different in the UK.⁴⁸ New, upper tail UK companies do brilliantly at attracting venture capital finance. And the UK has a large and liquid corporate bond market, totalling around £500 billion for investment grade securities, which allows larger and better-established companies to raise money at long maturities in capital markets.⁴⁹

Where the UK financial system looks weaker is when serving smaller and less well-established companies, the type of which are more likely to occupy the lower tail. Most are not of sufficient scale or standing to attract capital market financing. For them, banks are a more natural source of financing. But options for bank financing appear to be more limited, certainly than in Germany with its three-tier structure.

Lending to the corporate sector by UK banks, at 6% of their assets, is around one third of the equivalent by German banks (Table 6). In relation to GDP, bank financing of companies is around half that in Germany. The UK's national development bank (the British Business Bank) has assets that are a small fraction of its German counterpart (KfW) (Table 7). There are plans to increase the capacity of the British Business Bank, but that would still leave it significantly smaller than KfW relative to GDP.

⁴⁸ Bank of England (2016) discusses the measurement and financing of "productive investment" in the UK.

⁴⁹ Belsham and Rattan (2017).

In sum, the UK's long tail problem is largely a diffusion rather than innovation problem. And this problem seems to have its roots in transfer barriers – barriers to transferring technology, know-how, people and financing – from the UK's thriving hubs to its striving spokes. Stronger, longer spokes are needed to reach the long tail.

The Bank of England and Monetary Policy

Since its inception, the Bank of England's role has been to secure stability in prices and the financial system. Historical experience illustrates these are essential ingredients for investment and innovation by companies and for rising productivity, pay and living standards in the economy.⁵⁰ By achieving its statutory objectives of price and financial stability, the Bank of England provides one of the necessary foundations for productivity.

It is also well-known, including from historical experience, that productivity depends importantly on a number of structural features of the economy, including levels of education and skills in the workforce, the quality and quantity of infrastructure and innovation and the scale of financing to companies. Price and financial stability are necessary conditions for rising productivity. But they are far from being sufficient ones.

When it comes to those structural features of the economy, central banks do not have the tools to affect lasting change. Central bank tools are cyclical, rather than structural, because their impact on the economy is temporary, not permanent. We do not build schools, colleges, houses, roads, railways or banks. Nor do we finance them. Those tools, rightly, are in the hands either of governments or private companies. So too is the financing of them.

Through their role in securing stability in prices and the financial system, central banks do have a clear interest in understanding the forces shaping productivity and the supply-side of the economy. Central banks can also help in diagnosing productivity problems and in identifying policies that could lift barriers to productivity improvement. That is the purpose of this lecture. It is also the reason several Bank colleagues have given recent speeches on this topic.⁵¹

Before turning to that, let me say something about monetary policy. The job of securing price stability falls to the Bank of England's Monetary Policy Committee (MPC), of which I am a member. The MPC helps to determine monetary conditions in the UK, with a view to hitting a 2% inflation target on a sustainable basis.

Monetary conditions in the UK remain highly accommodative by any historical metric. Bank Rate currently stands at 0.5%, as it has stood for much of the past decade. That is, by some margin, lower than at any time in the Bank's 324-year history, which it celebrates next month. With headline inflation running at around 2.4%, this means inflation-adjusted interest rates in the UK are currently significantly *negative*.

⁵⁰ As discussed, for example, in Carney (2017).

⁵¹ For example, Ramsden (2018) and Tenreyro (2018).

The MPC has also undertaken asset purchases amounting to almost £½ trillion since 2009. This has provided additional monetary stimulus to the UK economy, at its peak equivalent to a further reduction in interest rates of up to around 2.5 percentage points.⁵² Taking account of that would put inflation-adjusted interest rates into even more deeply negative territory.

These extraordinary monetary policy measures have been an important factor supporting growth in the UK since the global financial crisis. Without these measures, we estimate the economy would have been around 8% smaller and unemployment 4 percentage points higher.⁵³ Highly accommodative monetary conditions have helped maintain UK growth at steady, if unspectacular, rates for the past six years.

Recent data suggest UK growth continues to be steady, at or around its trend rate of roughly 1.5% per year. At the same time, the labour market has continued to perform strongly. Almost 3 million new jobs have been created over the past six years. The unemployment rate in the UK, at 4.2%, has fallen to its lowest levels in over 40 years. Measures of tightness in the labour market have picked up notably. And the Bank's estimates suggest little or no slack remains in the economy.

Consistent with that, and despite undulations in shorter-run measures, we have seen a steady pick-up in annual rates of wage inflation over the past 18 months, as well as in pay settlements. Private sector regular pay growth has risen from around 2% this time last year to just shy of 3% currently. And average pay settlements of around 2% last year have picked up to around 2.5% this. Surveys of pay and recruitment difficulties, including from the Bank's Agents, have gone up a gear comparing this year with last.

The rise in whole-economy pay has been held back in recent years by the 1% cap on public sector pay. This, arithmetically, has lowered whole-economy pay growth. To my mind, it is likely also to have had a shadowing effect on pay-setting behaviour in the private sector, by setting a very low norm or floor for wage bargaining among private companies, in what has been a tightening labour market.

Either way, that public sector pay cap is now being lifted and decisively so. In excess of 1 million NHS workers will receive an across-the-board pay rise of 3% this year. Some NHS workers, especially those in lower salary ranges, may receive pay rises in excess of 10%. That is a nice way for those workers to be celebrating the NHS's 70th birthday.

The lifting of the public pay cap for the NHS could well have knock-on effects, not only to pay settlements elsewhere in the public sector, but to private sector settlements too given the tightness of the labour market.

⁵² Estimated based on simple ready-reckoners. There is a high degree of uncertainty around these estimates. In the May 2014 *Inflation Report*, the MPC judged that QE had a peak effect on GDP of around 2.5% in 2013. The range of QE estimates in Joyce *et al* (2011) would suggest that would have a percentage point impact on annual inflation of roughly half that amount. Joyce *et al* (2011) also suggest as a ready reckoner that a 100bp cut in interest rates would boost inflation by 0.5pp.

⁵³ Bunn, Pugh and Yeates (2018).

This will add to the gradual, but now clear, upwards impetus to wage and cost pressures in the UK economy, at a time when headline inflation is already somewhat above the MPC's 2% target.

Earlier this week, I was in Wales on a regional visit. The clear message from a number of employers in a number of industries was that wage pressures had picked-up this year. In a tight labour market with mounting recruitment difficulties, companies said they now needed to pay up to retain as well as attract staff. This is a message echoed elsewhere across the Bank's Agency network, covering the whole of the UK.

If so, this is a significant development. Until recently, the pick-up in wage pressures had been confined to workers moving jobs. Those remaining in post had seen no pick-up in pay. They had been the missing link in the pay chain. In a tight labour market, that link now appears to be being forged. If workers who stick as well as twist are now getting pay rises, this will add impetus to cost and inflationary pressures.

It is against this background that the MPC signalled in its May *Inflation Report* that some further modest tightening of monetary policy was likely to be required over the next few years to return inflation sustainably to target, provided the economy evolved broadly in line with its projections. Any such tightening, the MPC emphasised, was likely to be gradual in pace and limited in extent.

At its meeting in June, the MPC voted to maintain Bank Rate at 0.5%. With two other MPC members, however, I voted to raise Bank Rate by 25 basis points. The aim in doing so was exactly in line with the May *Inflation Report* – to reduce modestly the degree of monetary policy accommodation, given a tight labour market and gradually mounting pay pressures, to return inflation to target sustainably.

Some context is important here. Voting for a 25 basis point rate rise, a full decade after monetary policy was first placed on an emergency setting, is hardly either surprising or radical. A Bank Rate rise of 25 basis points would still leave monetary conditions in the UK extraordinarily accommodative by any historical metric. And the aim in doing so is to *lower* the risk of needing to tighten policy *less gradually* in future and cause a *sharper* adjustment in the economy.

Truth be told, I would have voted to raise Bank Rate at the MPC's May meeting had data on the economy held firm. What we saw ahead of that meeting was a string of weak data suggesting consumer spending might be faltering. I believed there was option value in waiting to see if these data signalled the start of a lasting retrenchment by households, or were instead a temporary snow or statistical blip. With only a modest policy tightening needed over a number of years to return inflation to target, there was "no rush".

In the event, data on the consumer since the May MPC meeting has, virtually without exception, bounced back strongly. That includes measures of retail spending, consumer confidence and consumer credit. The underlying picture now appears to be one of gently rising household spending. This is being supported by highly accommodative credit conditions and now-positive growth in inflation-adjusted wages.

And then, of course, there is the World Cup. Without wishing to tempt fate, England's recent sporting success on the football field (and cricket pitch) has probably added to that feel-good factor among England-supporting consumers. The "smile count" on my recent visits to Wales and Scotland was also as high as I can remember, although I suspect that may have been the weather rather than the football.

Of course, there are always some data – and, indeed, some football teams – which disappoint to the downside. That is what data does. But waiting for something to turn up is not a prudent strategy in life. And waiting for everything to turn up is certainly not a prudent strategy for monetary policy. That is why the MPC has signalled, and I have voted for, a modest reduction in monetary stimulus provided the economy evolves as expected.

Monetary policy guidance like this is not unconditional; it depends, as it should, on how the economy performs. From my conversations with household and businesses, such guidance is both understood and helpful to them when planning their spending and borrowing. It is the direction of travel and likely destination of interest rates, rather than the precise timing of any move, that matters for their decisions and hence for the fortunes of the economy.

Creating Institutional Spokes

Let me conclude by discussing briefly one or two areas where it might be possible to make progress in lifting the barriers to productivity growth. These suggestions very much go with the grain of proposals in the Government's excellent white paper on industrial policy. Indeed, one strand of that plan involves investigating the scope for improving performance among the long tail of under-performing UK companies.

These proposals complement various excellent private sector initiatives underway. These include Jurgen Maier's *Made Smarter* proposals, which aim to improve digitization across UK businesses. As the earlier evidence on penetration rates of different technologies made clear, this is huge scope for improvement among UK companies here. A second initiative is Sir Charlie Mayfield and Tony Danker's *Be the Business* movement. Their important and exciting initiatives are focussed on lifting the long tail.

One structural solution to the long tail problem I sometimes hear mooted is simply to lop it off. This, it could be argued, solves the UK's productivity problem at source and at a stroke. For economic historians, this could be thought of as an act of Schumpeterian creative destruction.⁵⁴ For horror film-watchers, it could be thought of as an act of systematic zombie-slaying.

⁵⁴ Schumpeter (1942).

I believe such an approach is probably wrong-headed. Most tail companies are not zombies, overburdened by an insurmountable mountain of debt or a broken business model. Rather they are companies surviving, but not yet thriving. And they account for fully 80-90% of all jobs. They are not the tail; they are the dog. If you restricted creative destruction to true zombie – those whose productivity was negative – this would make little arithmetic difference to average productivity, raising it by perhaps 1%.⁵⁵

A more creative solution to the long tail problem would be to build a stronger diffusion infrastructure – that is to say, creating some stronger and longer diffusion spokes to attach to the UK's strong innovation hub. Let me discuss three approaches for creating those institutional spokes. All three are already underway, to some extent. Once in place, they could speed-up the process of technological trickle-down to the long tail.

(a) Supporting Supply Chains

Recent years have seen a lengthening and deepening of supply chains, nationally and internationally.⁵⁶ The largest UK companies have direct supply chain links with several thousand other companies, large and small, national and international. As well as benefitting efficiency and lowering costs, these supply chains provide a natural infrastructure that could be harnessed to spread good practice across companies.

Some companies already actively support firms in their supply chain to help improve their productivity, the fruits of which can then be shared. A set of structured projects are underway nationally, co-ordinated by *Be the Business*, which aim to use the supply chain infrastructure to disseminate knowledge and best practice across companies in the long tail.

As another example, I recently visited a company which had invested in a new technology centre. This enabled them to simulate all of the processes and products in their complex business. They had created, in effect, a “digital twin” of their business. This technology allowed them to evaluate the impact of potentially productivity-improving changes to their business model, cheaply and quickly. Indeed, it was being used in just this way to drive through those changes.

Once established, this digital twinning technology can be applied to any business, large or small, simple or complex. Indeed the company I visited was doing just that, inviting in companies from their supply chain to help them understand and improve their business productivity-wise, in a way which benefitted both sides. If that digital twinning model could be introduced across greater numbers of companies and lengthier supply chains, it could serve as a very effective diffusion engine for innovation.

One way of doing so, recently proposed by academic and inventor Philip Bond, would be to create a central digital-twinning infrastructure that could be available to any company across the UK. You could think of this

⁵⁵ Assumes workers with negative productivity flow into firms with zero productivity.

⁵⁶ Baldwin (2016).

as a new piece of critical national infrastructure – a publically-available digital platform for diffusing new processes and practices across the UK company sector, enhancing technological trickle-down to the furthest reaches of the long tail.

(b) Supporting Technology Transfer

A second area where progress could be made in creating an improved diffusion infrastructure is by harnessing the UK's university network. The UK's top-ranked universities are already playing a crucial role in establishing the UK as an international innovation hub. Their business parks serve as crucibles of creativity. They are an important reason why the UK ranks so highly on innovation and start-ups.

The question, raised in a speech of mine last month, is whether the UK's other universities could serve as spokes for these hubs, helping diffuse innovation more broadly across sectors and regions. In principle, using the university network in this way has advantages. Unlike the Catapults, these institutional spokes do not need to be created anew. And the universities already have extensive regional and sectoral reach.

Clearly, this would require a fairly significant repurposing of universities. Any new activities would also need adequate financing. But neither of these problems seems insurmountable. And looking ahead, it seems likely universities will need to rethink their purpose and focus, in the light of profound shifts in the future of work and the skills demands of students.⁵⁷

(c) Supporting Human Capital

Ideas percolate through people. The UK's labour market currently operates in a way which does not support the trickledown of knowledge as much as it might. If talent only circulates around existing productivity pools, an important diffusion channel is being blocked. Rather than smoothing out differences, worker movements appear to strengthen the magnetic pull of the productivity poles, the upper tail North and the lower tail South.

The question is whether that magnetic pull could be reversed, or at least reduced, by supporting skills and staff transfers between the two poles. One way of doing so is through company mentoring or twinning. In effect, this is doing for companies what is often already being done for their staff, matching companies with very different levels of skills and experience.

Be the Business is already engaged in a sequence of pilot mentoring schemes between companies, enabling the sharing of skills and knowledge. This is about to become a nationwide mentoring programme. They are also developing programmes to develop management and technical skills within companies – “mini-MBAs” – to help close some of the skills gaps with other countries.

⁵⁷ Haldane (2018).

Beyond that, it may be worth considering whether a formal network for knowledge-exchange, along the lines of the German *Steinbeis* system, is worth considering. This would provide a more formal, systematic, institutional means of enabling skills transfer between companies, sectors and regions. This, too, would support diffusion and productivity growth in the long tail.

Conclusion

The UK faces perhaps no greater challenge, economically and socially, than its productivity challenge. Meeting that challenge would deliver benefits to workers in improved wages and skills and to companies in greater efficiency and profitability. It would also contribute to closing inequalities of income, wealth and opportunity which have rightly and increasingly pre-occupied policymakers over recent years.

The UK has a rich, in some respects world-leading, endowment of innovation and talent. This is, however, unevenly spread. Developing an institutional infrastructure, which draws on the UK's comparative advantage in innovation but which spreads its benefits more widely, would support the long tail of UK companies and the people who work for them. It would help close the pay and productivity gaps between the best and the rest, the present and the past, the in-crowd and the out. It would put the rhyme back into R&D.

The returns to doing so are difficult to quantify precisely. As a thought-experiment, imagine the bottom three quartiles of the UK productivity distribution saw their productivity gap with the quartile above closed. That would boost UK levels of productivity by around 13%.⁵⁸ This would close a large part of the productivity shortfall relative to its pre-crisis trend. And it would make inroads into closing the productivity gap with the US and Germany. In today's prices, it would boost the level of UK GDP by around £270 billion.

In closing those gaps, a useful intermediate objective would be to create in the UK a leading-edge diffusion infrastructure, to rival and complement its leading-edge innovation infrastructure. This boost our world (and, with luck, our World Cup) rankings. Inclusive innovation could serve as a conduit to inclusive growth. The UK's innovation hub would get the spokes it needs to reach every sector, every region, every worker. It would be an industrial strategy for everyone.

⁵⁸ Haldane (2017).

References

Acemoglu, D (2002), 'Directed Technical Change', *The Review of Economic Studies*, Vol. 69, No. 4, pp. 781-809.

Acemoglu, D, Gancia, G and Zilibotti, F (2015), 'Offshoring and Directed Technical Change', *American Economic Journal: Macroeconomics*, Vol. 7, No. 3, pp. 84-122.

Acemoglu, D and Robinson, J (2012), *Why Nations Fail: The Origins of Power, Prosperity and Poverty*, Penguin Random House.

Aghion, P, T, Besley, J, Browne, D, Caselli, R, Lambert, R, Lomax, C, Pissarides, N, Stern, J, van Reenen (2017), 'Investing for Prosperity: skills, infrastructure and innovation', *LSE Growth Commission Report*.

Andrews, D, Criscuolo, C and Gal, P (2015), 'Frontier firms, technology diffusion and public policy: micro evidence from OECD', *OECD Productivity Working Paper Series*, No 2.

Andrews, D, Criscuolo, C and Gal, P (2016), 'The Global Productivity Slowdown, Technology Divergence and Public Policy: A Firm Level Perspective', *Hutchins Center Working Papers*.

Baddeley, M (2006), 'Convergence or Divergence? The Impacts of Globalisation on Growth and Inequality in Less Developed Countries', *International Review of Applied Economics*, Vol. 3, pp. 391-410.

Bahaj, S, Foulis, A and Pinter, G (2017), 'Home values and firm behaviour', *Bank of England Staff Working Paper*, No. 679.

Baldwin, R (2016), *The Great Convergence: Information Technology and the New Globalization*, Harvard University Press.

Bank of England (2016), 'Understanding and measuring finance for productive investment', *Bank of England Discussion Paper*, April 2016.

Batten, S and Jacobs, D (2017), 'Foreign-owned firms and productivity', *Bank Underground*, 17 August 2017.

Belsham, T and Rattan, A (2017), 'Corporate Bond Purchase Scheme: design, operation and impact', *Bank of England Quarterly Bulletin*, Vol. 57, No. 3, pp. 170-181.

Berlingieri, G, Blanchenay, P and Criscuolo, C (2017), 'The Great Divergence(s)', *OECD Science, Technology and Innovation Policy Papers*, No. 39, May 2017.

Bjelland, M, Fallick, B, Haltiwanger, J and McEntarfer, E (2011), 'Employer-to-Employer Flows in the United States: Estimates Using Linked Employer-Employee Data', *Journal of Business & Economic Statistics*, Vol. 29, No. 4.

Bloom N and van Reenen, J (2007), 'Why Do Management Practices Differ Across Firms and Countries?', *Journal of Economic Perspectives*, Vol. 24, No. 1, pp. 203-224.

Brennan, S, Haldane, A and Madouros, V (2010), 'The Contribution of the Financial Sector: Miracle or Mirage?', in Turner, A *et al*, *The Future of Finance: The LSE Report*, London School of Economics and Political Science.

Brynjolfsson, E and McAfee, A (2014), *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W. W. Norton & Company.

Bunn, P, Pugh, A and Yeates, C (2018), 'The distributional impact of monetary policy easing in the UK between 2008 and 2014', *Bank of England Staff Working Paper*, No. 720.

- Carney, M (2017)**, opening remarks to the Bank of England 'Independence – 20 years on' Conference, available at <https://www.bankofengland.co.uk/speech/2017/opening-remarks-at-the-boe-independence-20-years-on-conference>
- Caselli, F and Coleman II, WJ (2006)**, 'The World Technology Frontier', *American Economic Review*, Vol. 96, No. 3, pp. 499-522.
- Coase, R (1960)**, 'The Problem of Social Cost', *The Journal of Law & Economics*, Vol. 3, pp. 1-44.
- Comin, D and Mestieri, M (2013)**, 'Technology Diffusion: Measurement, Causes and Consequences', *NBER Working Paper*, No. 19052.
- Comin, D, Trumbull, G and Yang, K (2011)**, 'Fraunhofer: Innovation in Germany', *Harvard Business School*, Case 9-711-022, March 2011. Revised 2012.
- Cornell University, INSEAD and WIPO (2017)**, *The Global Innovation Index 2017: Innovation Feeding the World*, Ithaca, Fontainebleau, and Geneva.
- Davis, S and Haltiwanger, J (2014)**, 'Labor Market Fluidity and Economic Performance', *NBER Working Paper*, No. 20479.
- De Loecker, J and Eeckhout, J (2017)**, 'The Rise of Market Power and The Macroeconomic Implications', *NBER Working Paper*, No. 23687.
- Department for Business, Energy & Industrial Strategy (2017)**, 'Industrial Strategy: building a Britain fit for the future', *UK Government Policy Paper*.
- Ernst & Young (2017)**, 'UK Catapult Network Review', *Ernst & Young*, SBS PS17086.
- European Commission (2018)**, *European Innovation Scoreboard 2018*, European Commission.
- Ferguson, N (2018)**, *The Square and the Tower: Networks, Hierarchies and the Struggle for Global Power*, Penguin.
- Feyrer, J (2009b)**, 'Trade and Income – Exploiting Time Series in Geography', *NBER Working Paper*, No. 14910.
- Feyrer, J (2009b)**, 'Distance, Trade, And Income – The 1967 To 1975 Closing Of The Suez Canal As A Natural Experiment', *NBER Working Paper*, No. 15557.
- Filippetti, A and Archibugi, D (2011)**, 'Innovation in times for crisis: National Systems of Innovation, structure and demand', *Research Policy*, Vol. 40, No. 2, pp. 179-192.
- Frankel, J and Romer, D (1999)**, 'Does Trade Cause Growth?', *American Economic Review*, Vol. 89, No. 3, pp. 379-399.
- Gordon, R (2012)**, 'Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds', *NBER Working Paper*, No. 18315.
- Haldane (2009)**, 'Rethinking the Financial Network', speech available at <https://www.bankofengland.co.uk/-/media/boe/files/paper/2009/banking-on-the-state>
- Haldane, A (2017)**, 'Productivity Puzzles', speech available at <https://www.bankofengland.co.uk/speech/2017/productivity-puzzles>
- Haldane, A (2018)**, 'Ideas and Institutions – A Growth Story', speech available at <https://www.bankofengland.co.uk/speech/2018/andy-haldane-speech-given-at-the-oxford-guild-society>

Haskel, J and Westlake, S (2017), *Capitalism without Capital: The Rise of the Intangible Economy*, Princeton University Press.

HM Government (2017), 'Building our Industrial Strategy', *HM Government Green Paper*, January 2017.

Jacobs, M, Hatfield, I, King, L, Raikes, L and Stirling, A (2017), 'Industrial Strategy: Steering structural change in the UK economy', *IPPR Commission on Economic Justice*.

Jones, R (2016), 'Innovation, research and the UK's productivity crisis', *Sheffield Political Economy Research Institute*, Paper No. 28.

Joyce, M, Tong, M and Woods, R (2011), 'The United Kingdom's quantitative easing policy: design, operation and impact', *Bank of England Quarterly Bulletin*, Vol, 51, No. 3, pp. 200-212.

Khwaja, AI, Mian, A and Qamar, A (2011), 'Bank Credit and Business Networks', *HKS Faculty Research Working Paper Series*, RWO11-017, Harvard University, John F. Kennedy School of Government.

Macpherson, N (2018), 'Pay rises will not solve Britain's productivity problem', *Financial Times*, 23 March 2018.

Manyika, J, Lund, S, Bughin, J, Woetzel, J, Stamenov, K and Dhingra, D (2016), 'Digital Globalization: The New Era of Global Flows', *McKinsey Global Institute*, McKinsey & Company.

Monetary Policy Committee (2014), 'Inflation Report', *Bank of England*, May 2014.

Moykr, J (2014), 'Secular stagnation? Not in your life', *VoxEU*, 11 August 2014.

North, D (1990), *Institutions, Institutional Change and Economic Performance*, Cambridge University Press.

Ramsden, D (2018), 'The UK's productivity growth challenge', speech available at <https://www.bankofengland.co.uk/speech/2018/dave-ramsden-panel-remarks-cbi-event>

Remes, J, Manyika, J, Bughin, J, Woetzel, J, Mischke, J and Krishnan, J (2018), 'Solving the Productivity Puzzle: The Role of Demand and the Promise of Digitization', *McKinsey Global Institute*, February 2018.

Rhodes, C (2017), 'Business statistics', *House of Commons Library*, Briefing Paper, No. 06152.

Schneider, P (2018), 'The UK's productivity puzzle is in the top tail of the distribution', *Bank Underground*, 29 March 2018.

Schumpeter, J (1942), *Capitalism, Socialism and Democracy*, Harper & Brothers.

Swinney, P (2018), 'The wrong tail', *Centre for Cities*, 24 May 2018.

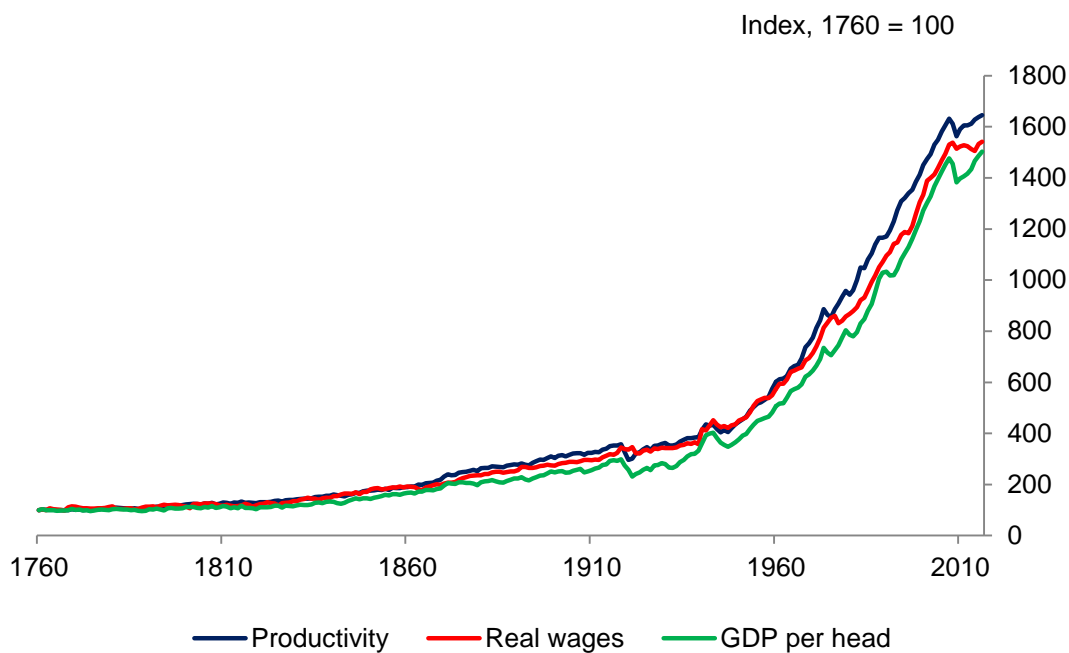
Tenreyro, S (2018), 'The fall in productivity growth: causes and implications', speech available at <https://www.bankofengland.co.uk/speech/2018/silvana-tenreyro-2018-peston-lecture>

Watts, D (2003), *Six Degrees: The Science of a Connected Age*, William Heinemann.

Wren-Lewis, S (2017), 'Disentangling the UK productivity problem', *mainly macro*, 28 November 2017.

Annex

Chart 1: UK productivity, real wages and GDP per head



Sources: Bank of England, 'A millennium of macroeconomic data', available online.

Notes: Labour productivity per head uses GDP at factor cost. Real wages defined as nominal wage divided by GDP deflator at factor cost.

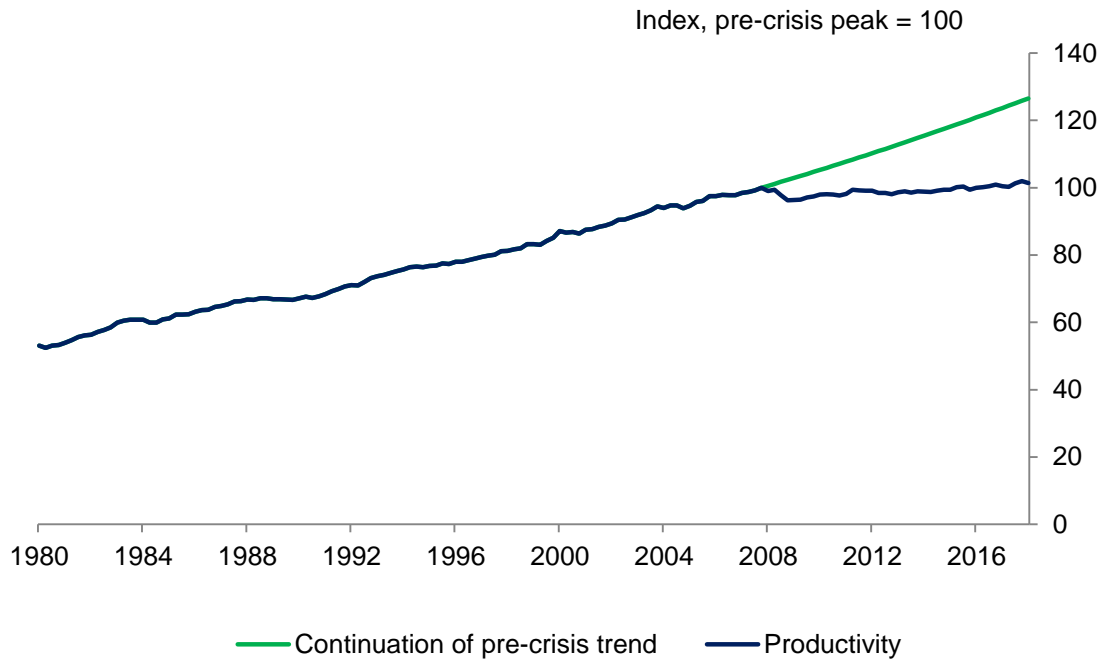
Chart 2: UK productivity and real wages since 1990



Sources: ONS and Bank calculations.

Notes: Productivity defined as output per head; real wages defined as whole economy average weekly earnings divided by consumer prices.

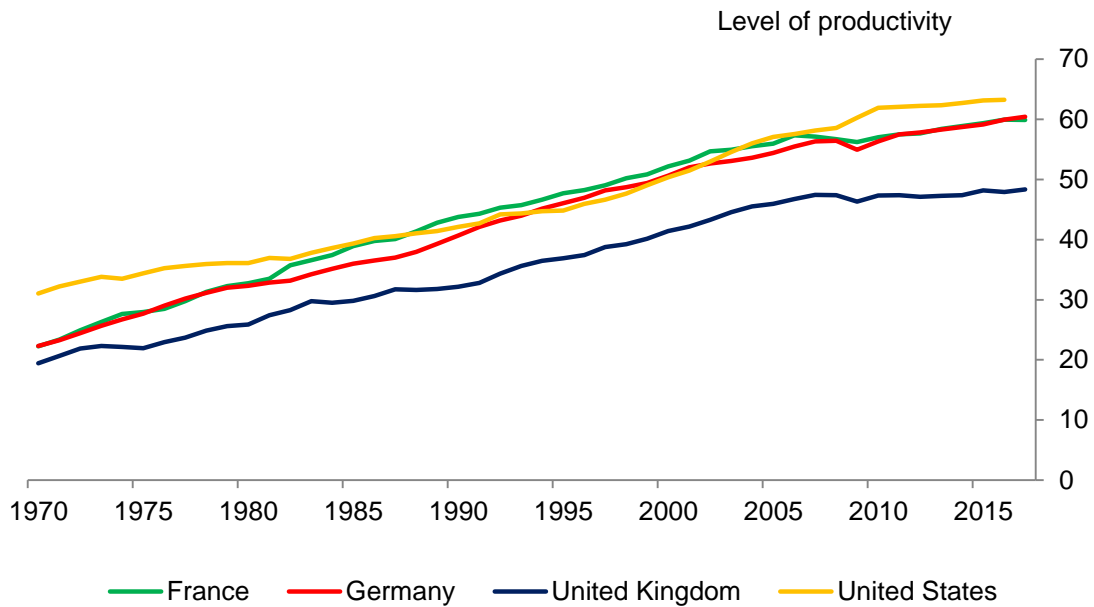
Chart 3: UK productivity relative to continuation of pre-crisis trend



Sources: ONS and Bank calculations.

Notes: Productivity is output per hour worked. Extrapolation of pre-crisis trend based on continuation of average growth from 1971 and 2007 beyond 2007Q4.

Chart 4: Productivity in UK, US, Germany and France




Sources: OECD and Bank calculations.

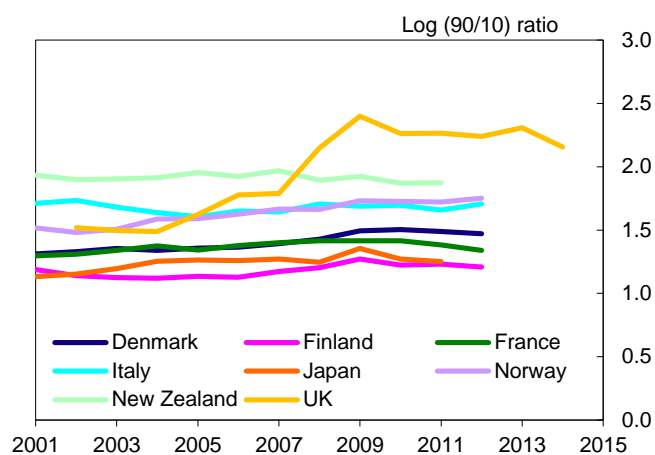
Notes: Productivity is output per hour worked in US\$ at constant prices, using 2010 PPPs.

Chart 5: Services and manufacturing firm-level productivity dispersion

Services (levels)



Manufacturing (levels)




Sources: OECD and Berlingieri, Blanchenay and Criscuolo (2017); ONS Research Database and Bank calculations.


Notes: Charts show the log difference between the 90th and 10th percentiles.

Chart 6: Change in firm-level productivity dispersion since 2001

Services (indexed)



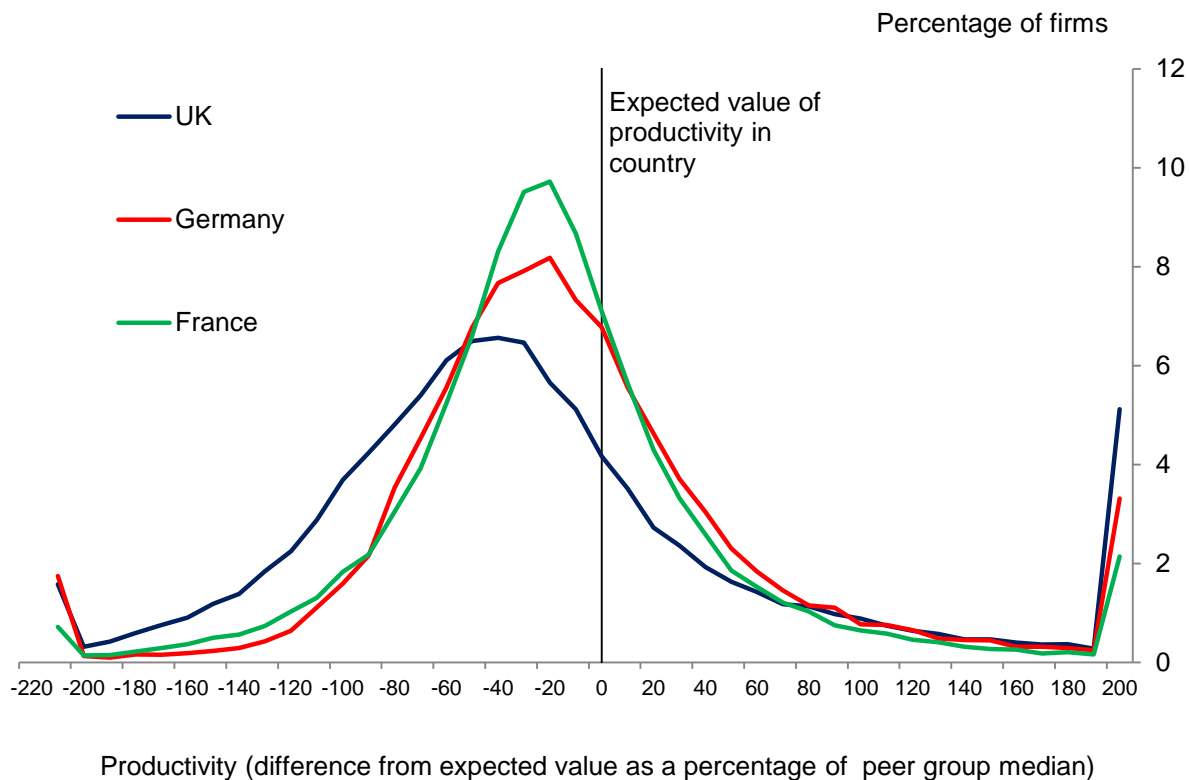
Manufacturing (indexed)



Sources: OECD and Berlingieri, Blanchenay and Criscuolo (2017); ONS Research Database and Bank calculations.

Notes: Chart show log difference between 90th and 10th percentiles indexed to zero in 2001. UK data only available from 2002, so UK base year = 2002.

Chart 7: UK, Germany and France firm-level productivity (data for 2013)



Sources: McKinsey and Orbis (2013).

Notes: Data kindly provided by McKinsey Global Institute. Estimated GVA (EBIT + employee costs) is regressed on a range of variables to control for sub-sector and number of employees using a Weighted Least Squares method (with employee numbers as the weighting). The output of this regression is used to compute an expected productivity, representing the average for a firm of that size in that sub-sector. The residual for each firm is plotted as a percentage of the median productivity for a firm in the same size bracket in the same sub-sector.

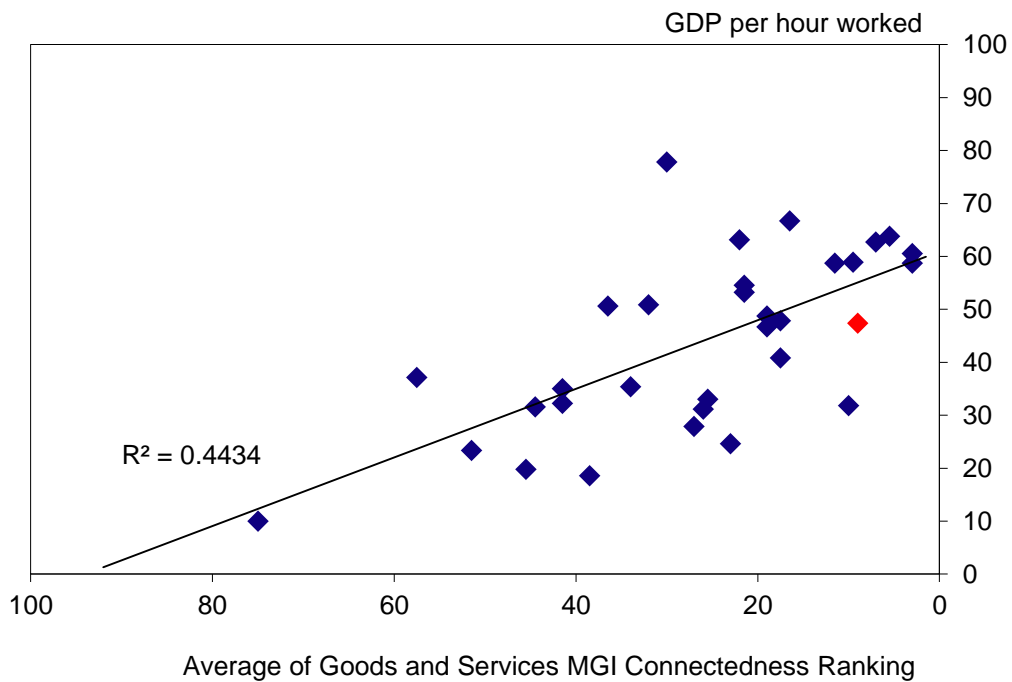
Table 1: Growth in the firm-level productivity distribution (aggregate, by region and by sector)

Annualised growth in average firm-productivity between 2004 and 2014			
	Top 0.1%	Top 1%	99% of firms
Aggregate	12%	8%	1%
Region			
North East	7%	5%	0%
North West	6%	4%	0%
Yorkshire & Humberside	12%	3%	-1%
East Midlands	30%	15%	1%
West Midlands	19%	10%	0%
East of England	16%	8%	0%
London	9%	8%	2%
South East	11%	6%	1%
South West	23%	12%	1%
Wales	0%	4%	1%
Scotland	13%	7%	0%
Sector			
Manufacturing	12%	7%	-1%
Electricity	-10%	-14%	-13%
Water	25%	8%	1%
Construction	14%	8%	-2%
Wholesale trade	4%	4%	0%
Transportation and storage	41%	28%	-1%
Information and communication	32%	18%	3%
Real estate activities	5%	3%	-1%
Professional, scientific and technical activities	9%	6%	4%
Administrative and support service activities	25%	14%	4%
Education	20%	13%	0%
Health	17%	7%	2%
Arts and entertainment	9%	4%	-5%
Other services	10%	0%	-5%

Sources: ONS Research Database and Bank calculations.

Note: Data refer to non-financial business sector. Calculations over the ten-year period in question are only possible for firms that exist for entire decade period and sampled throughout; these are likely to be larger firms.

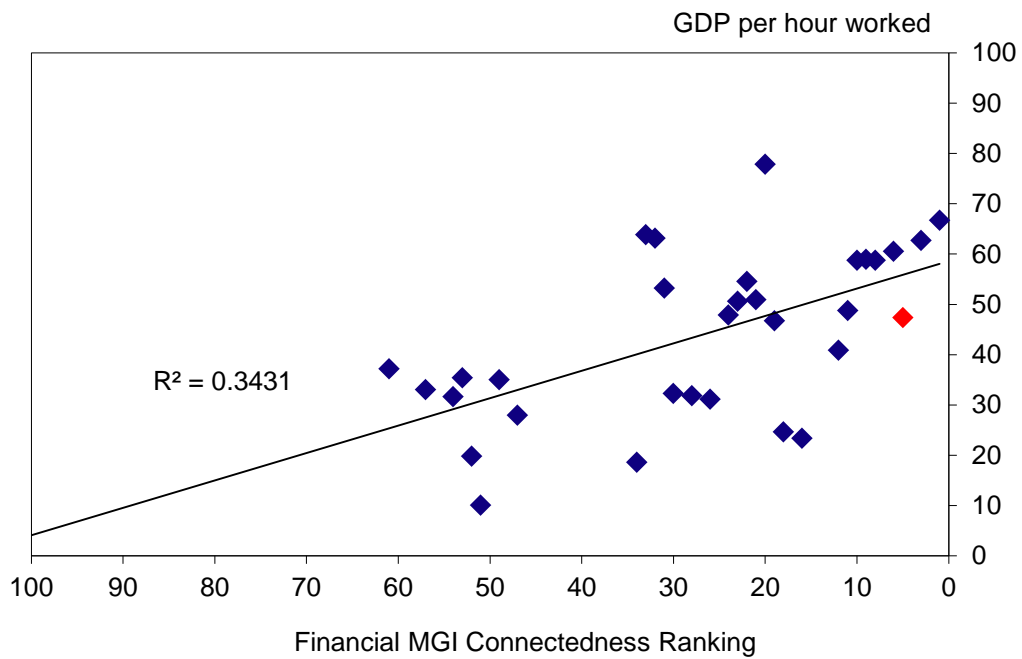
Chart 8: Trade openness and productivity



Sources: McKinsey Global Institute, OECD and Bank calculations.

Notes: Data for 2014. GDP per hour worked in US\$ at constant prices, in 2010 PPPs. UK marked in red.

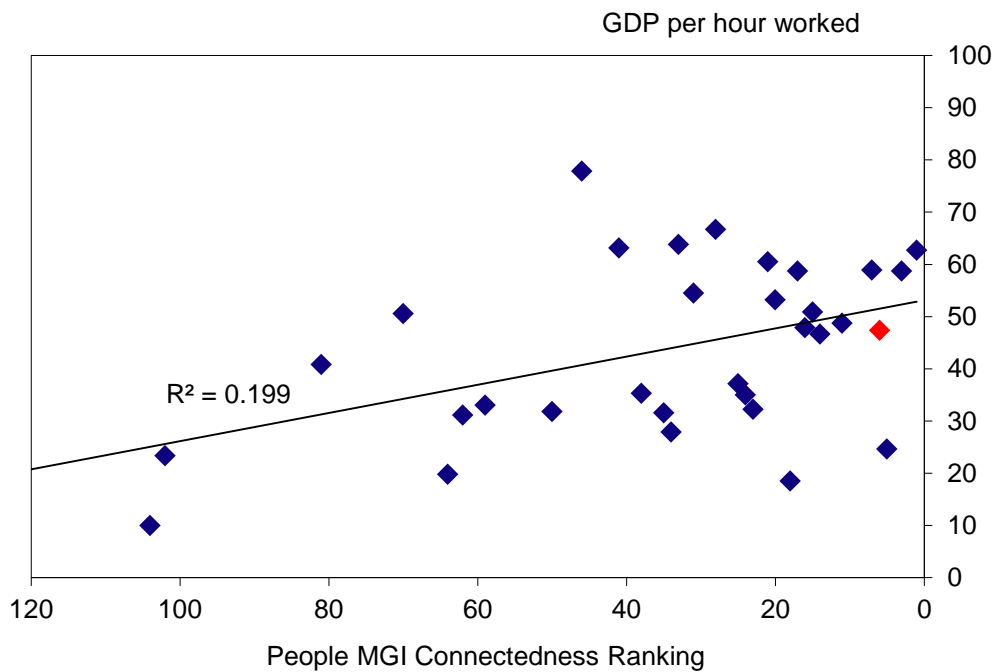
Chart 9: Finance openness and productivity



Sources: McKinsey Global Institute, OECD and Bank calculations.

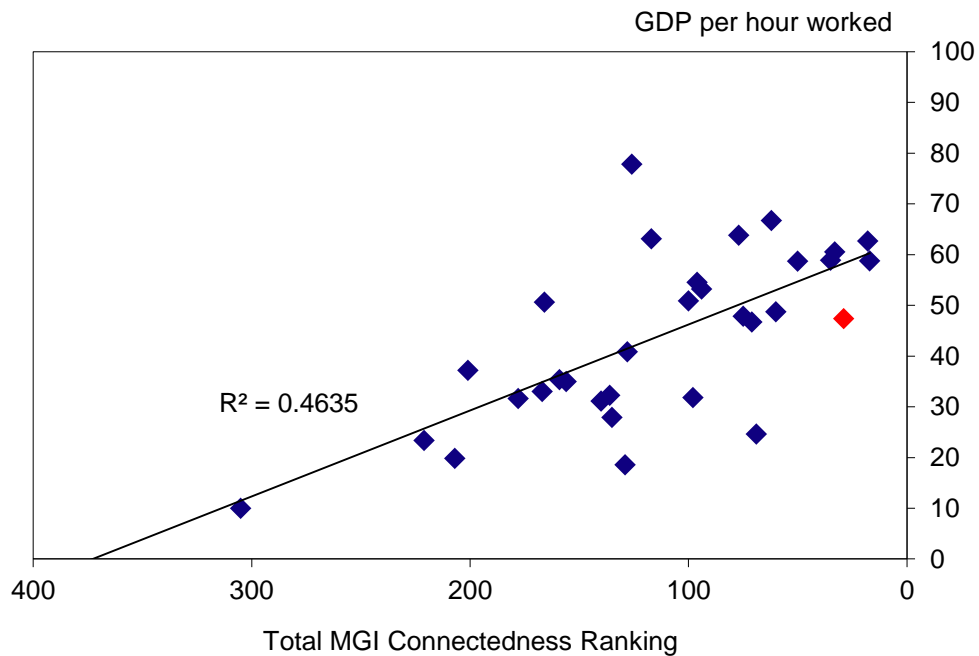
Notes: Data for 2014. GDP per hour worked in US\$ at constant prices, in 2010 PPPs. UK marked in red.

Chart 10: People openness and productivity



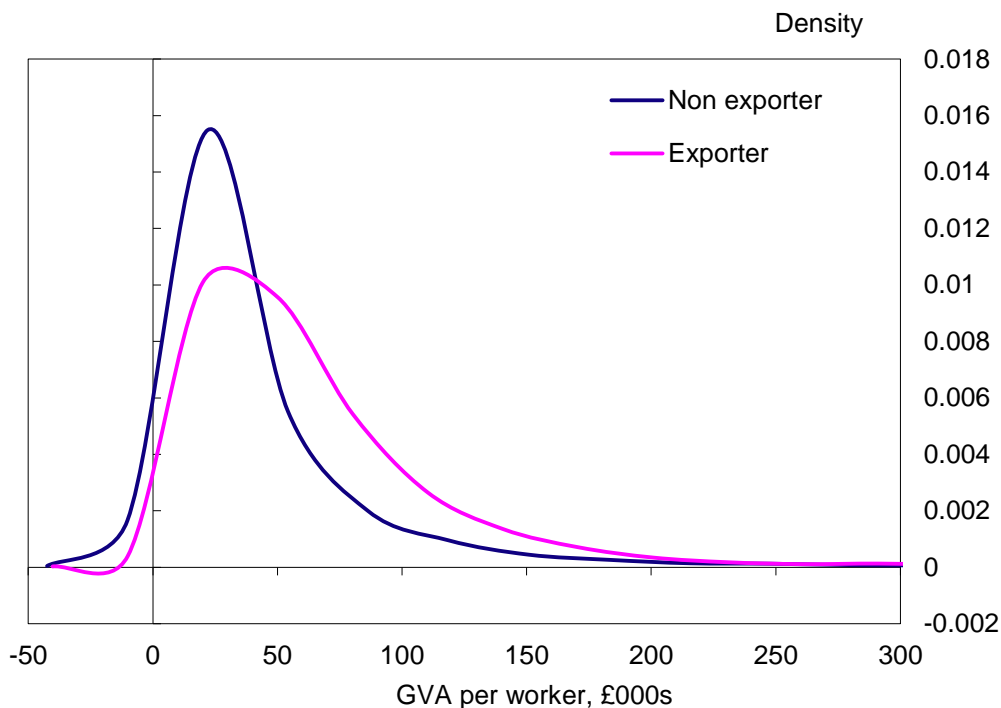
Sources: McKinsey Global Institute, OECD and Bank calculations.
Notes: Data for 2014. GDP per hour worked in US\$ at constant prices, in 2010 PPPs. UK marked in red.

Chart 11: Overall openness and productivity



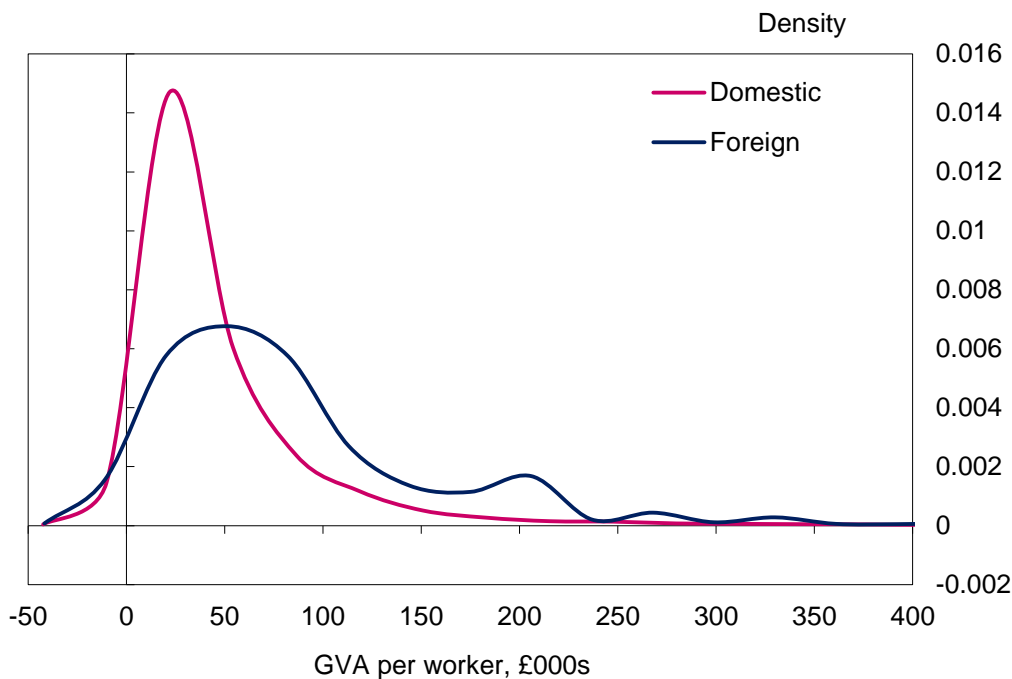
Sources: McKinsey Global Institute, OECD and Bank calculations.
Notes: Data for 2014. GDP per hour worked in US\$ at constant prices, in 2010 PPPs. UK marked in red.

Chart 12: Productivity for exporting and non-exporting firms



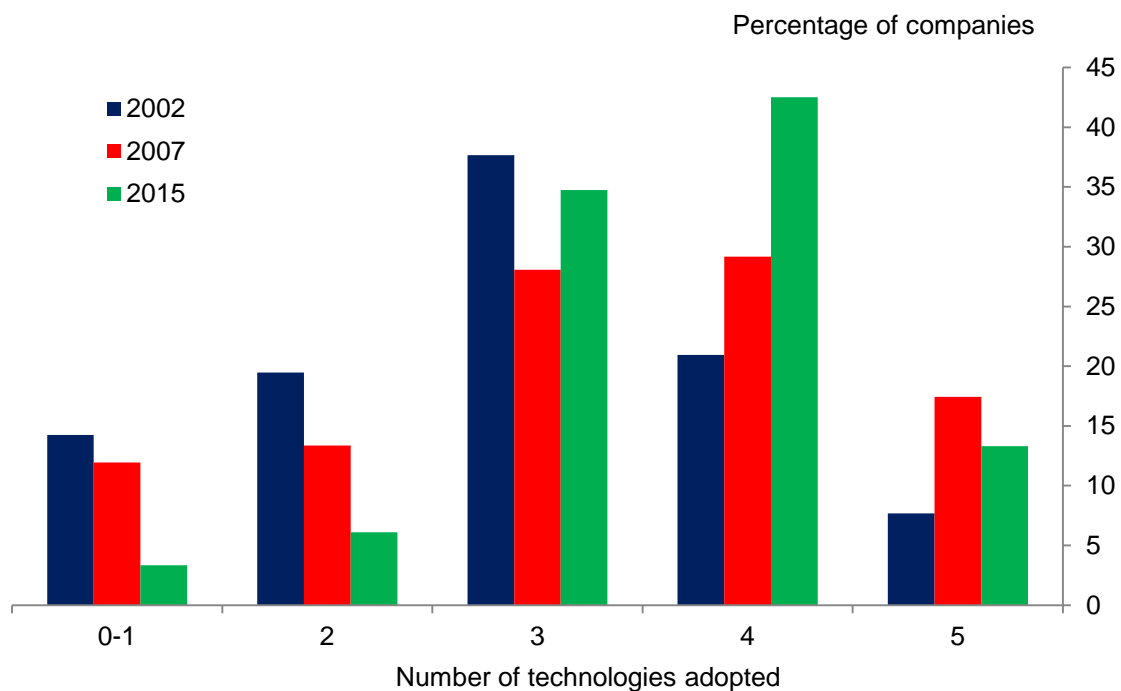
Sources: ONS Research Databases and Bank calculations.
 Notes: Kernel density estimate; real GVA per worker for 2013 and 2014.

Chart 13: Productivity for foreign-owned and domestically-owned firms



Sources: ONS Research Databases and Bank calculations.
 Notes: Kernel density estimate; real GVA per worker for 2013 and 2014.


Chart 14: Adoption rates of basic ICT technologies



Sources: ONS Research Databases and Bank calculations.

Notes: Basic ICT technologies refer to computers, internet, websites, e-purchases and e-sales.


Chart 15: Adoption rates of advanced ICT technologies



Sources: ONS Research Databases and Bank calculations.

Notes: Advanced ICT technologies refer to: mobile access to emails, documents and software, websites with online ordering, fast broadband access and electronic data interchange sales.

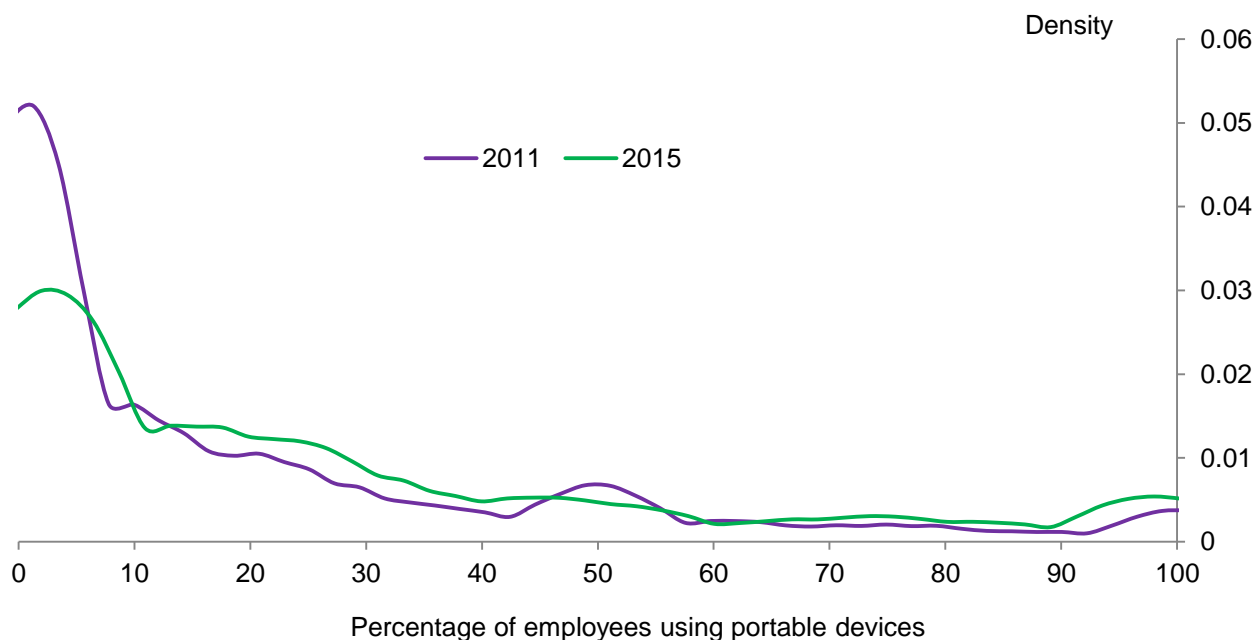
Chart 16: Employee usage of computers



Sources: ONS Research Databases and Bank calculations.

Notes: Chart shows responses to the question, "What percentage of people in this business use computers for their work?"

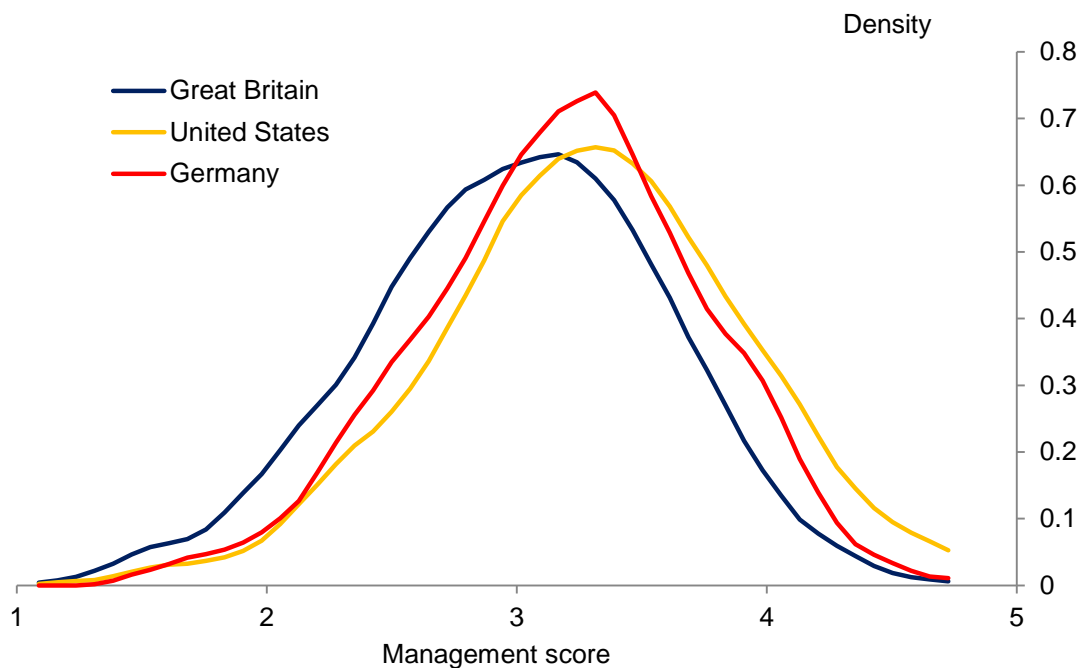
Chart 17: Employee access to a portable device with internet access



Sources: ONS Research Databases and Bank calculations.

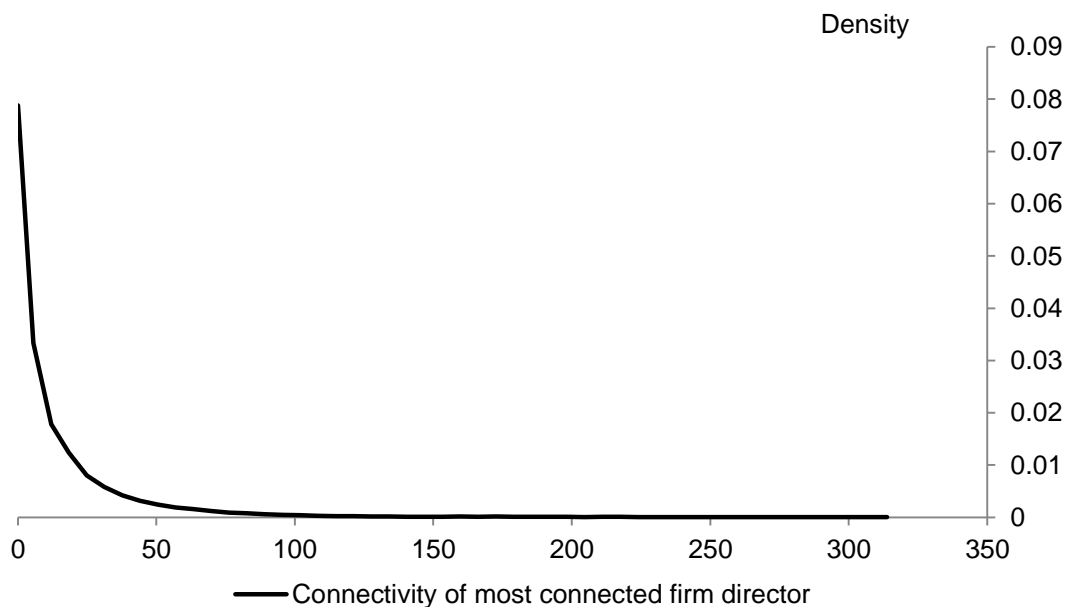
Notes: Chart shows responses to the question, "What percentage of people in this business use portable devices provided by this business that allow a mobile connection to the internet for business purposes?"

Chart 18: Distribution of management scores across countries



Sources: World Management Survey and Bank calculations.
 Notes: Kernel density estimate.

Chart 19: Distribution of director connectivity



Sources: Bureau van Dijk and Bank calculations.
 Notes: Measurement of director connectivity as in Khwaja, Mian and Qamar (2011). Connections of a director of a firm is defined as the number of co-directors the given director has at other firms. The connection measure takes value 0 when none of the directors of a given firm are directors of other firms.

Table 2: Distribution of director connectivity

	Percentiles								
	1%	5%	10%	25%	50%	75%	90%	95%	99%
Connections of most connected director via other firms	0	0	0	1	6	20	45	74	354
Number of directors at a firm	1	2	2	3	4	6	8	10	16

Sources: Bureau van Dijk and Bank calculations.

Notes: Measurement of director connectivity as in Khwaja, Mian and Qamar (2011).

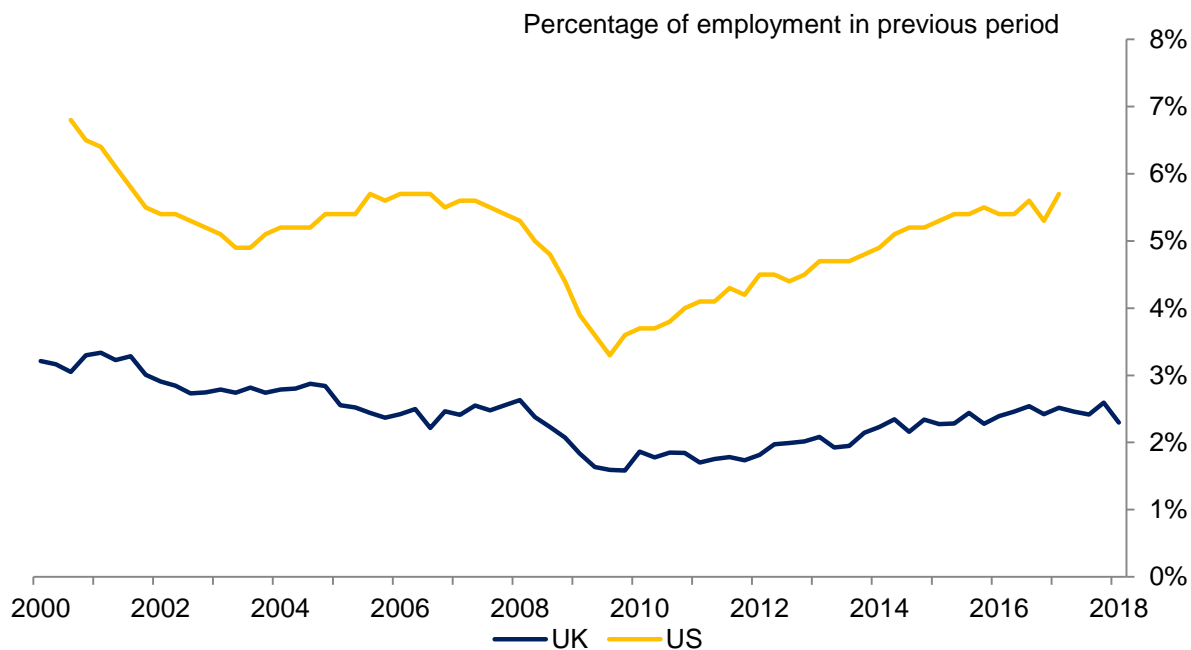
Table 3: Regression of firm-level productivity growth on director centrality

	$\Delta Labour Productivity_{i,t,t+1}$			
	(1)	(2)	(3)	(3)
Size	-0.1292*** (-19.28)	-0.1289*** (-19.05)	-0.1295*** (-19.34)	-0.0022* (-2.13)
$\log(NumberOfDirectors_{i,t})$	0.0161*** (3.66)		0.0138** (3.03)	0.0113*** (3.55)
$\log(Centrality_{i,t})$		0.0037*** (4.21)	0.0030*** (3.25)	0.0020*** (3.09)
Year FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	No
N	647787	647787	647787	673340
R^2	0.135	0.135	0.135	0.002

Sources: Bureau van Dijk and Bank calculations.

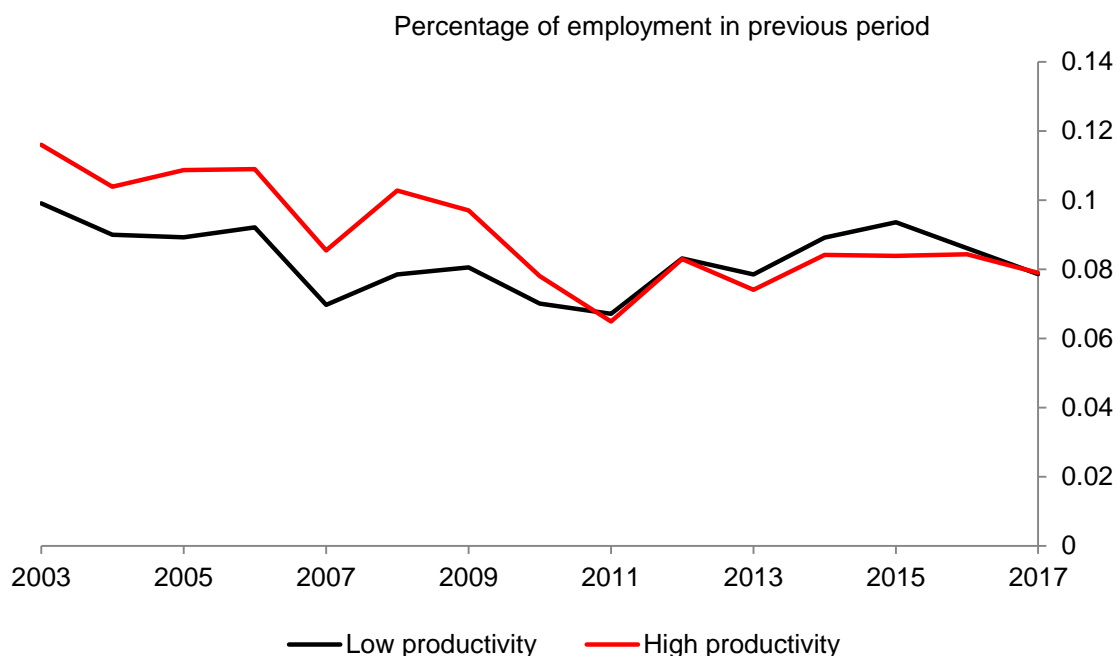
Notes: The table shows results from a panel regression of yearly productivity growth on firm size, the natural logarithm of the number of directors and the natural logarithm of firm centrality. Productivity is measured as gross value added per employee; size is measured as the natural logarithm of the balance sheet item "Total Assets"; centrality is measured as the number of connections (with other directors via other companies) of the most connected directors of the given firm. The sample period covers 2000-2013. T-statistics correspond to standard errors that are clustered at the firm-year level. Significance levels are * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Chart 20: UK vs. US labour market churn



Sources: ONS, United States Census Bureau and Bank calculations.
 Notes: Data show job-to-job transition rate as a share of employment.

Chart 21: Job churn by firm-productivity



Sources: ONS and Bank calculations.
 Notes: 'Low productivity' defined as the average churn rate per decile for those firms in the bottom half of the productivity distribution; 'high productivity' defined as the average churn rate per decile those firms in the top half of the productivity distribution.

Table 4: Job moves and productivity

		Decile at time <i>t</i>				Increasing productivity
		2 - 3	4 - 5	6 - 7	8 - 9	
Decile at <i>t-1</i>	2 - 3	0.41	0.26	0.14	0.18	Increasing productivity
	4 - 5	0.21	0.29	0.21	0.30	
	6 - 7	0.20	0.25	0.25	0.31	
	8 - 9	0.20	0.19	0.21	0.40	

→ Increasing productivity

Sources: ONS and Bank calculations

Notes: Each cell shows the probability of a worker being employed in a given quantile of the firm productivity distribution conditional on where in the distribution that worker was in the previous year, and having moved jobs.

Table 5: UK Catapults vs. German Fraunhofer

Dimension	UK Catapults	German Fraunhofer
Age	7 years (from 2011)	69 years (from 1949)
Number of centres	10	72
Industry projects per year	600	6,000 – 8,000
Annual budget	£0.3bn	€2.3 billion
Annual budget (% of nominal GDP)	0.01%	0.07%
Employees	> 1,400	24,500

Sources: catapult.org.uk, Ernst and Young (2017), Comin, Trumbull and Yang (2011) and fraunhofer.de/en

Notes: UK Catapults have delivered around 4,100 industry collaborations in total to date, which equates to around 600 per year. Their funding was £1.4bn over a five-year period, which equates to around £0.3bn per year.

Table 6: UK and German bank lending to businesses

Dimension	UK	Germany
% of total bank assets	6%	18%
% of nominal GDP	23%	49%

Sources: Bank of England, Deutsche Bundesbank, ONS, Thomson Reuters Datastream and Bank calculations.

Note: UK data refer to lending to non-financial corporates by UK resident monetary financial institutions (MFIs) for 2018Q1. German data refer to lending to enterprises and the self-employed by Germany resident Banks in 2018Q1.

Table 7: British Business Bank vs. KfW

Dimension	British Business Bank	KfW
Age	4 years (2014)	70 years (from 1948)
Total assets	£1.4bn	€472.3 billion
Assets as a share of annual nominal GDP	<0.1%	16%
Employees	144	6,284

Sources: British Business Bank, KfW Group, ONS, Eurostat and Bank calculations.