



The Banks that Said No: the Impact of Credit Supply on Productivity and Wages

Jeremy Franklin¹ · May Rostom² · Gregory Thwaites³

Received: 22 December 2017 / Revised: 4 January 2019 / Accepted: 17 January 2019 /
Published online: 03 April 2019
© The Author(s) 2019

Abstract

This paper estimates the effects of changes in bank credit supply on the real economy. We use UK firm-level data around the global financial crisis and information on pre-existing bank lending relationships to isolate exogenous credit supply shocks. We find some evidence that contractions in credit supply substantially reduce labour productivity, wages, and capital per worker within firms, and increase the chance firms will fail. Our results have implications for the welfare costs of financial crises, and for the costs of policy measures affecting credit supply at other times.

Keywords Firm behaviour · Productivity · Credit supply · Banks

JEL Classification D22 · D24 · G21

1 Introduction

The financial crisis of 2008 was associated with falls in the growth of corporate lending, business investment, labour productivity and real wages in many major industrialised countries. What were the causal links between these events? Did firms retrench because they could not get financing? Or did they become pessimistic about demand for their products, and demand both less financing and fewer factors of production as a result?

✉ Jeremy Franklin
jeremy.franklin@bankofengland.co.uk

May Rostom
may.rostom@bank-banque-canada.ca

Gregory Thwaites
gregory.thwaites@bankofengland.co.uk

¹ Bank of England, London, UK

² Bank of Canada, Ottawa, Canada

³ Bank of England and London School of Economics, London, UK

These questions are important for central banks in understanding the implications of financial crises for bank lending and how this affects the macroeconomy. The answers are important for assessing the costs and benefits of banking reforms intended to reduce the risk of financial crises.

This paper provides new evidence on the impact of credit supply on corporate outcomes. Our aim is to identify the impact of the reduction in credit supply following the 2007/8 financial crisis on labour productivity, investment behaviour and average pay, as a guide to the potential costs of changes in credit supply in other circumstances. We employ a new identification strategy that exploits information on pre-crisis lending relationships within a large firm-level dataset of non-financial companies. We use the UK - a large, bank-dependent industrialised economy that suffered a very serious banking crisis in 2008-9 - as a lens through which we can examine a question that is generic to any modern economy.

All firms in the UK are required to register the identity of any party (a 'chargeholder') that has a claim on the firm's assets as collateral for a loan. We construct a proxy for pre-crisis banking relationships by identifying major banks among these chargeholders. We show that these relationships are persistent, and that they help to predict the amount firms borrow after the crisis. We exploit the stickiness of these relationships, together with the fact that different banks tightened credit conditions to different degrees, to generate exogenous variation in credit supply at the firm level. We then use this instrument for credit supply to quantify the impact of a change in total debt on firm outcomes, controlling for demand conditions in the product market.

We find that firms who faced a reduction in credit supply experienced larger falls in labour productivity, capital per worker and average pay. On our preferred specification, a 10% contraction in borrowing caused by credit supply led, on average, to a 5–6% fall in capital per head, a 5–8% fall in labour productivity, and a 7–9% fall in average pay for the affected firms. On the face of it, the reduction in credit supply may therefore be responsible for a large part of the fall in labour productivity growth observed in the UK, and potentially other advanced economies, since the crisis.

Our results are consistent with a mechanism whereby an increase in the shadow price of capital caused firms to substitute towards more labour-intensive forms of production, which in turn lowered labour productivity growth. The estimated impact on labour productivity is large, suggesting that lower credit supply may also have been associated with lower levels of innovation and technological development. Average pay also fell further in firms more exposed to the credit shock, and in similar proportion to labour productivity, even though these firms were hiring labour in the same markets as less exposed firms. This observation lends support to rent- or risk-sharing theories of wage determination (Van Reenen (1996)) - in other words, firms were able to share some of their idiosyncratic productivity shocks with their workers. We also find that firms facing adverse credit supply shocks were more likely to fail.

Relative to the existing literature, our study makes two principal contributions. First, our paper supports existing evidence on the impact of credit supply on macro outcomes. In particular, it explores the links to productivity in an economy which saw a large fall in productivity after the crisis. Second, our paper is the first to use bank relationships to study the credit shock in the United Kingdom, an economy which is both heavily dependent on banks and which suffered a relatively large credit shock.

The remainder of this paper is structured as follows. In Section 2 we discuss how our work relates to existing studies of the impact of credit on corporate outcomes. Section 3

briefly reviews relevant aspects of the UK banking system, corporate sector and macroeconomic situation. Section 4 presents the dataset used in our analysis. Section 5 sets out our empirical approach and identification strategy. Section 6 presents our results and a number of robustness tests. Section 7 discusses the economic significance of our results. Section 8 concludes.

2 Existing literature

There are a number of studies that examine the importance of bank shocks on the provision of the supply of loans.

Kashyap and Stein (2000) show that the transmission of monetary policy is stronger for banks with less liquid balance sheets and that this affects their lending behaviour. While they do not discuss the economic impact of this contraction in lending, others have argued that this behaviour can lead to a slower growth. Peek and Rosengren (1997) find that Japanese banks cut lending in the US following deterioration in their parent banks' capital positions. They then go on to show in a follow up paper using regional data that this in turn affected US construction activity (Peek and Rosengren 2000). Gilchrist and Zakrajsek (2012) use aggregate data to identify the impact of a credit shock on the macroeconomy. They find that shocks to the excess bond premium leads to elevated risk aversion in the financial sector, and in turn to a contraction in lending supply and economic activity.

A key challenge for research in this area is to disentangle firms' demand for credit from banks' supply of loans. Microdata can help address this concern. Khwaja and Mian (2008) use matched firm-bank data to quantify the impact of a bank liquidity shock on the provision of loans to firms. More recently, others have taken this a step further: using matched firm-bank data to track the effect on overall economic activity. In particular, several papers have used exposure to different lenders just before the 2008 financial crisis, as a means of generating cross-sectional variation in credit supply during the crisis.

For the US, Greenstone and Mas (2012) use geographic variation in the pre-crisis market share of different banks across the US, along with variation in the credit crunch across banks, and finds that US counties in which poorly performing banks had bigger market shares saw fewer new loans, less employment and fewer business start-ups during the crisis. Edgerton (2012) uses data on lending relationships for a sample of equipment finance loans to identify the impact of restricting supply of credit to firms. He finds that variation across lenders accounted for around 17% of the decline in aggregate equipment financing, or about one-third of the total decline in financing of small businesses. Flannery et al. (2013) find that US firms which had relationships with banks with higher non-performing real estate loans borrowed less and invested less following the crisis. Chodorow-Reich (2014) measures banking relationships by identifying the lead arrangers for a given firm's syndicated loans. Having found that bank-firm pairs are sticky, he uses 'distressed' lenders as a proxy for restricted credit supply. He finds that employment fell more sharply during the crisis among the clients of less healthy lenders, particularly when those clients were small firms. The withdrawal of credit can explain roughly one-third of the employment decline in the sample in the year following the Lehman collapse.

Outside of the US, Amiti and Weinstein (2013) use matched Japanese bank-firm data over the period 1990-2010 to decompose loan movements into bank, firm, industry and common shocks. They find that idiosyncratic bank shocks have a large effect on investment.

Paravisini et al. (2015) estimate the elasticity of exports to credit using matched Peruvian customs and firm-level bank credit data. They compare changes in exports of the same product and to the same destination in order to account for non-credit determinants of exports. They then compare the outcomes of firms borrowing from different banks that were differently affected by the 2008 financial crisis. Their results suggest that the reduction in credit reduced exports by raising the variable cost of production. Bentolila et al. (2013) merge the Spanish credit register with balance sheet data and find that Spanish firms who entered the crisis with relationships to weak banks experienced larger falls in employment.

Ongena et al. (2013) examine how corporate outcomes of firms that are dependent on credit differed from those that are not credit constrained. They focus on firms located in Eastern Europe and Asia, since the region was not initially affected by the global financial crisis. Their identification strategy relies on distinguishing between 3 types of banks according to whether they are domestic, foreign-owned or able to borrow on international wholesale markets. They find that banks with access to international wholesale funding cut back their lending by more than domestic banks; and that firms dependent on credit from these banks had lower returns on asset growth and lower revenue growth.

In summary, there is a small but growing literature using bank relationships to study the effect of credit supply on corporate outcomes, principally borrowing, investment and employment. Our paper is the first to do so in the UK, and to look at labour productivity, capital investment and average pay.

3 Macroeconomic context

3.1 Corporate access to credit in the UK

Firms in the United Kingdom are highly dependent on banks as a source of debt finance. The top six banks account for 70% of the stock of lending to UK firms (Bank of England 2013). In fact, only about 250 firms have access to the public bond market and they account for only 12% of private employment. In addition, it is typically only larger firms that have access to equity markets with only a very small fraction of small firms gaining access.¹

In their reliance on banks, instead of bond and equity markets, UK firms are much closer to continental European firms than US firms; bank loans account for about three-quarters of euro area corporate debt, about two-thirds in the UK, and about one-quarter in the US Pattani et al. (2011).

Broadly speaking, the reason for firms' reliance on bank lending is because the cost of getting a rating from a credit reference agency is too high for many small firms – a prerequisite required by investors who rely on these scores as a signalling device about the riskiness and viability of the firm. In other words, variation in the credit supplied by banks is, for most UK firms, coterminous with variation in the overall level of credit available to them.

¹A 2012 Department of Business Innovation and Skills (BIS) report found that 'a minority [of SMEs] use equity finance' and the proportions are 'too small to show on a graph' (van der Schans 2012). Instead, they find that half of all SME's have used financial institutions to obtain finance. The remaining half did not seek funding or used other means not listed in their survey. Overall, BIS estimates suggest that in 2014 SMEs, which make up 97% of all firms in the UK, raised only £1.1 billion from private external equity compared to nearly £40 billion from bank lending British Business Bank (2014).

3.2 UK banks and the financial crisis

The major UK banks - Barclays, HSBC, Lloyds Banking Group and the Royal Bank of Scotland (RBS) - had very different experiences during the recent financial crisis. Lloyds TSB and HBoS (Halifax Bank of Scotland) merged to form the Lloyds Banking Group (LBG). LBG, together with RBS, were subsequently part-nationalised by the UK government after a £50 billion capital injection in October 2008. Figure 1 shows the premia on credit default swaps on the senior unsecured bonds of the big four UK banks, a measure that is highly correlated with their funding costs. These were very similar and at record lows before the crisis, but became high and dispersed after.

A necessary condition for our identification strategy is that this dispersion was not in large part caused by systematic differences in the health of UK banks' corporate loan books. Official narrative accounts of the failures of HBoS and RBS (FSA 2011; PCBS 2013) support this idea, laying the blame instead on trading book losses and reliance on wholesale funding. The key exception to this were the large losses made in RBS' and HBoS' commercial real estate (CRE) portfolios. For this reason, and in line with Bentolila et al. (2013), we exclude CRE firms from our sample. Once excluded, we find no obvious correlation between the change in banks funding costs and their exposure to lending in the corporate sector. (For more detail on the relationship between bank's funding costs and their exposure to CRE and non-CRE firms, please refer to Appendix A).

3.2.1 Credit supply following the financial crisis

By 2008, UK banks were under intense pressure to recapitalise. To help alleviate funding pressure, some initiatives were set up by the government and the Bank of England (e.g. the Special Liquidity Scheme (SLS) and the Credit Guarantee Scheme (CGS) in April and October 2008 respectively). However, banks themselves had started to issue equity in order to improve their capital position. The top 4 UK banks issued £58.9 billion and £60 billion in equity in 2008 and 2009 respectively, but issued none in the years before and after that. Because the banks suffered shocks for different underlying reasons, the timing and quantities issued varied by bank. For example, HSBC issued no equity in 2008, but did so in 2009; Barclays issued equity in 2008 but not in 2009; and RBS issued equity in both years.²

As a result, banks started to cut back on lending, in particular during the acute period of 2008 and 2009. By 2008, lending growth to UK non-financial corporations had slowed down dramatically. Data from the Bank of England suggest that annual (3-month on 3-month) growth rate in corporate lending fell by 20 percentage points between 2007 and 2008, after it had been growing at an average rate of roughly 10% a year in the previous decade.

By 2010, the UK banking system had stabilized and funding positions had improved (although they still needed to build capital to move towards full Basel III compliance). Most of this was done organically by seeking to increase the proportion of their books funded by deposits.

This narrative chimes with the results presented in Section 5. We find that the effects are strongest for 2008 and 2009 (although the instrument is weaker in 2009). Our preliminary investigations found that by 2010 the results start to fade. While this may be due to firms switching banks or ceasing to trade, it is also consistent with the timing of capital issuance outlined above.

²This is important for our identifying assumption as not all the banks failed and when they did suffer serious problems, these were not occurring in unison.




Fig. 1 UK bank CDS spreads. Source: MarkIT; authors calculations

3.3 The productivity puzzle

The UK experienced a deep and prolonged economic recession following the 2008 financial crisis. Labour productivity fell significantly in the aftermath of the crisis, and has stagnated since. This weakness in productivity has been puzzling, and is associated with surprisingly strong employment rather than weak output. It has also been weaker in the UK than in many other advanced economies. A detailed discussion of the UK's productivity performance since the crisis is provided in Barnett et al. (2014a).

The weakness in labour productivity has also coincided with a sharp fall in UK corporate borrowing and real wages (Fig. 2). A key outstanding question, and the focus of this paper, is the extent to which the disruption in credit supply witnessed in the aftermath of the 2008 crisis has been a cause of the weakness in investment, productivity and wages.

This post-crisis weakness in productivity has not just been confined to the UK, but is increasingly becoming a global phenomenon (Fig. 3). The UK provides an interesting case study into this issue since, like many other European economies, it has a corporate sector that is heavily reliant on banking as its main source of credit.




Fig. 2 UK macroeconomic data, relative to trend. The figure shows the percentage point difference of each data series relative to its pre-crisis trend. Apart from GDP, all measures refer to the private sector. Real PNFC bank credit refers to the stock of private non-financial corporation (PNFC) loans, excluding the commercial real estate (CRE) sector, deflated using the GDP deflator. Trends are calculated as the log linear trend between 1993 and 2007 Q2. Source: ONS; authors calculations




Fig. 3 Labour productivity in the G7. The figure shows output per hour across G7 countries. Source: ONS; authors' calculations

4 Data

Our dataset is compiled from information taken from the Bureau Van Dijk (BvD) FAME database. This service extracts information from UK companies' annual accounts that are submitted to Companies House - the UK business registry for all limited companies. Companies House is responsible for incorporating and dissolving companies, and as such, all limited companies in the UK are required to register with Companies House. The BvD database contains information on around 1.2m registered UK companies.

4.1 Sample selection

The aim of this study is to examine the impact on companies of a contraction in the supply of credit from banks, measured through changes in companies' total borrowings. The key variable for our analysis is therefore the level of total debt held by individual companies, which we define as the total amount of overdrafts, short term loans and long term debt.

Since we are interested in examining the effects from the financial crisis, we select a cohort of firms that had positive levels of total debt in 2007. We then track the behaviour of these firms after the crisis.

The reporting requirements for Companies House vary by firm size.³ Large companies are required to send their full balance sheet and profit and loss accounts to Companies House; medium-sized companies can send abbreviated profit and loss accounts; and small companies can send abbreviated balance sheet data and are not required to send profit and loss information. As a result, there will be missing information on a number of smaller firms which are then dropped from our sample.

In common with Khwaja and Mian (2008) and Amiti and Weinstein (2013) and Chodorow-Reich (2014), the process of constructing a sample with all the variables we need reduces the sample size substantially and skews it towards larger firms. This means that there will be missing information on smaller firms, leading them to be underrepresented in

³In 2007, to be classified as a small (medium) company by Companies House, a firm had to meet at least two of the following criteria for two consecutive years: (i) annual revenue was less than £5.6 million (£22.8 million), (ii) the balance sheet total was less than £ 2.8 (£11.4) million, or (iii) there were less than 50 (250) employees.

Table 1 Total number of firms by size. Columns (1) and (2) are taken from 2007 ONS Business Population Estimates. Columns (3) and (4) show the number of firms in the final sample

	Population		Our sample	
	Number (1)	% (2)	Number (3)	% (4)
1-49	1,186,115	97.32	4,150	30.87
50-249	26,690	2.19	5,582	41.52
250+	5,920	0.49	3,712	27.61
Total	1,218,725	100.00	13,444	100.00

Source: ONS; author's calculations

our sample. Table 1 compares the number of small, medium and large firms in ONS estimates of the business population with our sample. Overall, our final sample contains over 10,000 firms per year, of which 72% are small and medium sized companies (SMEs) with less than 250 employees and 28% are large companies with over 250 employees.

4.2 Variable construction and summary statistics

Our identification strategy relies on information regarding pre-crisis relationships between companies and individual banks. To get this we extract information on registered charges from the BvD FAME database. A charge is the security a company gives for a loan and must be registered at Companies House within 21 days. The BvD FAME database captures information on persons entitled to an outstanding charge raised at Companies House. It also includes information on when the charge was created and when it ended (when the loan matured).

Table 2 compares summary statistics for the key variables in our sample of firms in 2007. Precise definitions are listed in Appendix B.⁴

A key variable in our analysis is firm-level labour productivity per head. There are two broad approaches to measuring gross value added (GVA) at the firm level, either (i) as the sum of profits and wages or (ii) as revenue minus non-labour intermediate costs. Both approaches have drawbacks. The two main issues with (i) are lower reporting rates and the presence of negative values (since firms reporting a loss could have negative GVA) which introduces difficulties when comparing percentage changes over time. The main issue with (ii) is how best to measure non-labour costs, which are not well reported in annual accounts. In order to maximise coverage in our sample, we instead choose to proxy GVA with revenue. For firms where both GVA and revenue could be calculated, the correlation between estimates of GVA per head and revenue (or turnover) per head was very high.⁵ For the remainder of this paper, results for labour productivity refer to revenue per head.

⁴As is typically the case in firm-level data, the distributions of all our variables are heavily skewed to the right. We also compare an equivalently sized random sample of firms that do not report information on total debt. This second sample of firms is on average much smaller, both in terms of total capital and employment. They also have much lower reporting rates across the set of variables, suggesting that a large proportion are exempt from reporting more detailed financial information on account of their size.

⁵A simple regression of the change in revenue per head on the change in productivity, controlling for 2 digit industries, yielded a coefficient of 0.8 for 2008 and 0.9 for 2009 and both were statistically significant at the 1% level.

Table 2 Summary statistics, 2007. Total debt is the total amount of long-term debt and short-term overdrafts. Capital per worker is the total amount of tangible assets divided by the number of employees. Revenue per head is the total revenue divided by number of employees. Average pay is total remuneration divided by number of employees. Employment is the number of employees. Standard deviations are shown in square brackets

Variable	Mean	Median	Min	Max	Obs
Total debt (£'000s)	78,880 [1,019,833]	3,213	0	92,000,000	13,444
Capital per worker (£'000s)	317 [4,835]	17	0	408,474	13,444
Revenue per head (£'000s)	315 [2,503]	122	0	230,680	13,444
Average pay (£'000s)	34 [36]	27	0	2,006	13,444
Employment	924 [8,116]	102	1	507,480	13,444

5 Empirical approach and identification strategy

5.1 The baseline regression specification

Our aim is to quantify the effect of a credit supply shock on various aspects of firm performance and behaviour, principally labour productivity, capital per worker, wages and firm survival.

The first question to address is how to quantify the size of the shock itself. If a firm suffers a negative credit supply shock, the amount it borrows will tend to fall. One natural and convenient choice of metric is therefore the amount of debt a firm has borrowed. Our baseline equation is thus

$$\Delta y_{it} = \beta_{0t} + \beta_{1t} \Delta d_{it} + \beta_{2t} x_{it} + \epsilon_{it} \quad (1)$$

where i indexes firms, t is time, d is the stock of a firm's debt, y is the response variable of interest and x is a vector of observable firm characteristics. The vector of observables x is chosen with a view to making the residual variation in pre-crisis firm characteristics uncorrelated with our instrumental variable, and comprises the following variables: 2-digit industry sector; whether the firm is a subsidiary, a parent company or a standalone firm; the log level of revenue in 2007, and firm age in 2007 at time t .^{6,7}

The amount a firm borrows will be driven by both the supply of and demand for credit. For example, a firm might reduce borrowing because of a reduction in credit supply, but also because it might want to dispose of physical capital or otherwise alter its capital structure. In

⁶We add a value of 1 to the level of revenue in order to include firms with zero revenue.

⁷When firms cease to operate, the left-hand side variable is typically not recorded or recorded at zero such that the log difference is undefined. We use symmetric growth rates - as is now common in this literature - to measure Δy_{it} and Δd_{it} . For instance, taking a variable z , the change in z from 2007 to t is calculated as follows

$$\Delta z_t = \frac{z_t - z_{2007}}{0.5(z_t + z_{2007})}$$

each case, the correlation between credit and investment will be different. So a simple OLS regression of, say, investment on the change in debt will typically deliver biased estimates of the effect of a credit supply shock. Appendix C describes the determinants of the bias in OLS estimates of β , and shows that it can be of either sign.

To avoid this bias, our regression specification is the following two-stage least squares model:

$$\Delta d_{it} = \theta_{0t} + \theta_{1t}b_{i,2007} + \theta_{2t}x_{it} + \mu_{it} \quad (2)$$

$$\Delta y_{it} = \beta_{0t} + \beta_{1t}\Delta d_{it} + \beta_{2t}x_{it} + \epsilon_{it} \quad (3)$$

where $b_{i,2007}$ is a vector of seven indicator variables for the identity of the bank with which firm i had relationship with before the crisis.

The first four indicator variables represent a firm having a relationship with exactly one of the big four banks - Barclays, HSBC, Lloyds (including HBoS) and RBS (including NatWest). In view of the relative infrequency with which other banks appear in our sample, the remaining three respectively code for a relationship with a bank outside the big four, and a relationship with a non-bank, or a relationship for which there is no assigned chargeholder.

We exclude firms that have multiple banking relationships. These firms are different from our main sample as they tend to be much larger and include multinational corporations. These firms are likely to have a greater range of external funding sources available, making it easier to switch between sources of credit. These constitute only 6% of firms. Our IV results are qualitatively robust to reasonable alternatives to this scheme. For example, we get similar results when we separately code relationships with the big banks, irrespective of whether the firm has relationships with other banks.

One reasonable alternative to this specification would be to group banks according to whether they are strong or weak, and therefore more or less likely to provide credit, and then perform a difference-in-differences analysis comparing firms who have relationships with these two groups. The problem with this approach is that it is not obvious how to group banks. On one hand, banks like RBS and Lloyds became so weak that they were nationalised. On the other, nationalisation itself may have prompted a change in lending policy and actually boosted credit supply from the affected banks (see Rose and Wieladek (2014) for evidence that nationalisation affected the lending of UK banks). Furthermore, banks facing tougher funding conditions may have had more of an incentive to provide loan support or forbearance to weaker firms to avoid having to realise further losses on their balance sheet. This might have had a positive impact on measured lending. For example, latest estimates suggest the scale of forbearance across non-CRE SME borrowers is likely to have been relatively small by 2013, but could have been higher immediately after the crisis Arrowsmith et al. (2013).

A third approach could be to use an instrument of pre-crisis bank health or asset quality. We explored a few options here, however the information available in annual accounts pre-crisis limits the available options. We discuss how further work might help to address this in the final section.

Nevertheless, our two-stage least squares approach uses the variation across all banks to identify the impact of the credit shock. In Section 6.3.2 we rank and then group banks by the change in the volume of credit they supplied, and perform a difference-in-differences analysis on the resulting groups. We find evidence that corroborates our main story from this approach.

Equation 1 is essentially a time-differenced version of the specification in Paravisini et al. (2015). Our dataset is an unbalanced panel of firms, so we could in principle estimate




Fig. 4 New and outstanding loans. The figure shows the change in the proportion of firms with a loan from different institutions for the subsample of firms that had a loan to one of the big four banks in 2007. Source: BvD; authors' calculations

a variant of Eq. 1 in levels terms rather than first differences, controlling for time-invariant unobserved heterogeneity using standard methods. This approach would, however, suffer from a number of important problems. Most obviously, the impact of the credit shock may itself have varied over time. Furthermore, our identifying strategy, explained in detail below, relies on using pre-crisis banking relationships as an instrument for firm-level credit supply. Over the passage of time banking relationships will change and end for a variety of reasons. This means that the coefficients of Eq. 2 are highly likely to be unstable over time, and we verify that this is the case in the next section.

We therefore estimate our model in terms of changes between the year 2007, before the most serious phase of the credit crisis, and each of the post-crisis years 2008 and 2009 in our sample. Beyond 2009, our identification strategy begins to fail as pre-crisis banking relationships have longer to decay, and do so non-randomly, so we do not present results for later years.

5.2 Identification

Our identifying assumption is that a firm's banking relationships are correlated with its performance, conditional on observables, but only through the effect that bank identity has on credit supply. In terms of Eqs. 2 and 3, these assumptions are respectively

$$E[\Delta d_i b_i] \neq 0 \quad (4)$$

$$E[\epsilon_{it} b_i] = 0 \quad (5)$$

Figure 4 highlights that banking relationships in our sample appear to be very sticky. It shows the change in the percentage of companies with a loan to one of the big four banks, a different bank, and a non-bank.⁸ The figure shows that over 10% of these firms no longer had a loan from one of the big four banks by 2011— but only around 1% had taken up a new loan with a different bank, and just under 2% had taken up a loan from a non-bank.

The stickiness of these banking relationships suggests that pre-crisis relationships will be a determinant of post-crisis borrowing opportunities. Section 6.1 below confirms statistically that pre-crisis relationships are very strong predictors of post-crisis borrowing.

⁸This is for the subset of companies with an outstanding loan with one of the big four banks in 2007.

Together, these pieces of evidence establish that Eq. 4 - the relevance of our instruments - holds.

On the question of instrument validity, Eq. 5, there are two principal reasons why our identifying assumption might fail. The first possibility is reverse causation - whether firm performance led to changes in credit supply rather than the other way around. If the firms who had relationships with a given bank performed systematically worse (conditional on observables) than others, say because that bank had selected riskier or less promising borrowers than others, causation could run from corporate performance to bank relationships - violating our identifying assumption.

This reverse causation seems unlikely in practice. Section 3 sets out narrative evidence that the main cause of variability in banks' performance after the crisis was not their corporate lending decisions (apart from in commercial real estate (CRE)). Given that we exclude CRE firms from our sample, we are omitting the major potential source of reverse causation from our sample.

The second possibility is selection on unobservables, whereby a firm's performance and the lending behaviour of its bank are influenced by a common, unobserved factor. For example, suppose a given bank took above-average risks on both sides of its balance sheet in the lead-up to the crisis. In the event of a system-wide financial shock, its lending would have contracted more than average on account of funding difficulties, and its borrowers may have performed less well than average because they were more exposed to the economic cycle. We would then observe a conditional correlation between performance and borrowing at the firm level, but it would not be causal - another violation of our identifying assumption.

We provide three pieces of evidence which suggest that selection on unobservables is unlikely to be present in our sample to a significant extent. First, there was no significant correlation, conditional on our other control variables, between bank relationships and our left-hand side variables before the crisis. If banks had been selecting firms based on their characteristics, such as level of riskiness, then we might expect to see statistically significant differences in our response variables y across banks conditional on observables characteristics *before* the financial crisis. We estimate the following equation for a series of outcome variables, denoted by y_{it} , in the period before the financial crisis. The outcome variables are gross profit, capital, employment, capital per head, revenue per head, and average pay and the equation is estimated using data from 2005 - 2006

$$\Delta y_{it} = \phi_0 + \phi_1 \mathit{bank}_{it} + \phi_2 x_{it} + t + e_{it}$$

where Δy_{it} is the symmetric growth rate between (2005 and 2006), bank_{it} are the same set of bank indicator variables for whether the firm had a charge with one of the big four banks, another financial institution or a non-financial institution in 2007. We test whether the coefficients on the vector ϕ_1 of bank dummies are jointly equal to zero. Table 3 reports p-values

Table 3 Joint F-test of significance of bank dummies. The table reports p-values from joint F-test of the bank indicator vector $\mathit{bank}_{i,2006-2005}$, where the dependent variable is the $\Delta y_{i,2006-2005}$

	$\Delta y_{i,2006-2005}$	P-value
Gross Profit		0.611
Capital		0.081
Employment		0.486
Capital per head		0.159
Average pay		0.215
Revenue per head		0.823

from the joint F-test. The p-values show that we cannot reject results for the coefficients jointly being equal to zero on all outcome variables (although capital per head is slightly significant at the 10% level).

Second, and relatedly, in Section 6.3.1 below we re-run our baseline two-stage least squares analysis but move the sample dates and selection rules two years back in time, attempting to predict response variables in 2006 and 2007 - before the most serious phase of the crisis - on the basis of banking relationships in place in 2005. We show that our instrument is not relevant before the crisis. The second stage regressions are also largely insignificant. This further supports the idea that there is unlikely to be common causation driving our results.

Third, the regression tables report standard tests of overidentifying restrictions, which are typically not rejected at standard significance levels. Taken together, this evidence suggests that our instrument is both valid and relevant.

6 Results

6.1 Baseline results

6.1.1 First stage regression

Table 4 sets out the results of the first stage regression (Eq. 2) of the change in credit volumes on our vector of identifiers for bank relationships $b_{i,2007}$. The first columns show results from an unbalanced panel of firms while the second two columns show results for the balanced panel of firms who have data available in both years.⁹ The identities of the banks themselves are anonymised.

The variation across our bank coefficients is relatively small suggesting that the variation in lending by banks may not be large— at least according to our data set. Nonetheless, our specification pass several standard tests described below.

The last two rows of the table show that the regression as a whole, as well as the joint test on θ_{1t} , is highly significant. The coefficients on $b_{i,2007}$ themselves are precisely estimated and in many cases bilaterally significantly different from each other. Coefficients on the other control variables are not reported for the sake of brevity.

We calculate standard errors for our baseline regressions using a heteroscedasticity-robust estimator. At first sight, it might be warranted to cluster our estimates at the bank level. However, our instrument is already a bank-level fixed effect, making this redundant. Our first-stage equation regresses the change in debt at the firm level on the bank relationship dummies and other controls. Any variation that is common at the bank level will be absorbed by the bank dummies.

For the second-stage equation, our identifying assumption is shown in Eq. 5. This means that any residual variation at the bank level, of the sort that would require clustering of standard errors, would amount to a violation of our identification restrictions. In other words, if our standard errors are biased by not clustering, then so are the parameter estimates themselves. We provide evidence elsewhere in favour of our identifying assumptions. Furthermore, we have calculated clustered standard errors (not shown) and confirmed that they

⁹Note that firms which cease operation are also included in our sample. They are recorded as having zero output

Table 4 First stage - baseline sample. The table shows results from regressions on the change in debt, $\Delta debt_{it}$, on a vector of bank dummies in 2007. Each column represents a cross-sectional regression for the change in time t relative to 2007, where t is (2008, 2009). Regressions control for industry sector, firm entity type, age and log level of revenue in 2007. Balanced are for firms present in both years

	Unbalanced		Balanced	
	2008 (1)	2009 (2)	2008 (3)	2009 (4)
Bank 1	-0.084*** [0.02]	-0.076*** [0.02]	-0.092*** [0.02]	-0.067*** [0.02]
Bank 2	-0.065*** [0.02]	-0.086*** [0.03]	-0.081*** [0.02]	-0.084*** [0.03]
Bank 3	-0.072*** [0.01]	-0.064*** [0.02]	-0.077*** [0.02]	-0.057*** [0.02]
Bank 4	-0.054*** [0.01]	-0.044** [0.02]	-0.057*** [0.01]	-0.043** [0.02]
Other financial institution	-0.042 [0.03]	-0.051 [0.03]	-0.075** [0.03]	-0.031 [0.04]
Non financial institution	-0.051*** [0.01]	-0.051*** [0.02]	-0.054*** [0.02]	-0.047** [0.02]
Controls	YES	YES	YES	YES
Observations	13,444	11,418	10,471	10,471
R-squared	0.007	0.008	0.007	0.009
F statistic	3.692	3.911	8.939	3.840
P-value of regression	0.000	0.000	0.000	0.000
P-value of all banks	0.000	0.001	0.000	0.007

Notes: *Significant at 10% level **significant at 5% level *** significant at 1% level. P-values in square brackets

are generally very close to simple heteroscedasticity-consistent estimates. Given the inaccuracy involved in calculating cluster-robust errors with so few clusters (see e.g. Cameron and Miller (2015)), there is a risk that clustering might reduce, rather than improve, the precision of our estimates of the standard errors. So we calculate standard errors based on a heteroscedastic but diagonal error covariance matrix.

6.1.2 Second stage regression

Tables 5, 6, 7, 8 and 9 set out our main (second stage) results from Eq. 3. Each table shows all possible permutations of the time period in question (changes from 2007-2008 and 2007-2009), whether the sample is unbalanced or balanced (whether we restrict

Table 5 Second stage - baseline sample: Capital per head. The table shows results from the second stage regression of the change in tangible assets per head, Δy_{it} , on the change in debt, $\Delta debt_{it}$, which is instrumented using a vector of bank dummies in 2007. Each column represents a cross-sectional regression for the change in time t relative to 2007, where t is (2008, 2009). Regressions control for industry sector, firm entity type, age and log level of revenue in 2007. Balanced are for firms present in both years. Firms that have become inactive or are in liquidation are included in the sample as having zero revenue. Failed firms are included as having zero revenue and zero assets

	Unbalanced				Balanced			
	2008		2009		2008		2009	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in debt	0.119*** [0.01]	0.528*** [0.14]	0.128*** [0.01]	0.584*** [0.20]	0.122*** [0.01]	0.611*** [0.14]	0.127*** [0.01]	0.472** [0.21]
Controls	YES	YES						
Observations	13,444	13,444	11,418	11,418	10,471	10,471	10,471	10,471
Hansen-Sargan Stat		5.590		10.810		3.687		9.868
Hansen-Sargan P-values		0.348		0.055		0.595		0.079
Kleinbergen-Paap		7.227		3.632		7.381		2.967
P-value of excl. instruments		0.000		0.001		0.000		0.007

Notes: *Significant at 10% level **significant at 5% level *** significant at 1% level. P-values in square brackets

attention to observations available in both years) and for for both OLS or IV estimation of Eq. 3. We report the elasticity, β_{1t} , on the change in debt, Δd_{it} . We also report the Sargan-Hansen test of overidentifying restrictions, the Kleinbergen-Paap F-Statistics of weak instruments and the P-value from the Angrist and Pischke F-test of excluded instruments from the first stage of the regression.

We first present results for capital per worker in Table 5. In both years, and for both the balanced and unbalanced sample, the OLS estimates of β_{1t} are much smaller than their IV counterparts (although both are also statistically significant). In view of the argumentation in Section 5, this attenuation bias from OLS is consistent with most changes in firm borrowing being due to credit demand unrelated to the firm's operating business - i.e firms managing their balance sheets. Turning to the IV estimates, the coefficient for the unbalanced and balanced samples are similar in magnitude; the effects vary somewhat by year and sample, but by less than one standard deviation. Taken together, the IV results suggest that a 10% fall in borrowing due to a credit supply shock leads to a 5-6% reduction in capital per head. This suggests that a credit supply shock has large effects on the capital intensity of production at the firm level.

The diagnostic tests reported with the regression suggest that any bias in this estimate is unlikely to be very large due to any issues with instrument validity and/or relevance, especially in 2008: the Sargan test of overidentifying restrictions (where the null is that our instruments are valid) is not rejected at standard levels for either year, while the F-statistic for the joint significance of the instruments suggests that the probability of a large weak instrument bias is low for 2008 (Stock and Yogo 2002).

Table 6 Second stage - baseline sample: Revenue per head. The table shows results from the second stage regression of the change in revenue per head, Δy_{it} , on the change in debt, $\Delta debt_{it}$, which is instrumented using a vector of bank dummies in 2007. Changes in debt and y are calculated using symmetric growth rates. Each column represents a cross-sectional regression for the change in time t relative to 2007, where t is (2008, 2009). Regressions control for industry sector, firm entity type, age and log level of revenue in 2007. Balanced are for firms present in both years. Firms that have become inactive or are in liquidation are included in the sample as having zero revenue. Failed firms are included as having zero revenue and zero assets

	Unbalanced				Balanced			
	2008		2009		2008		2009	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in debt	0.005 [0.01]	0.826*** [0.16]	0.003 [0.01]	0.533*** [0.17]	0.003 [0.01]	0.798*** [0.15]	0.009 [0.01]	0.603*** [0.20]
Controls	YES	YES	YES	YES	YES	YES	YES	YES
Observations	13,444	13,444	11,418	11,418	10,471	10,471	10,471	10,471
Hansen-Sargan Stat		4.033		8.429		5.089		8.855
Hansen-Sargan P-values		0.545		0.134		0.405		0.115
Kleinbergen-Paap		7.227		3.632		7.381		2.967
P-value of excl. instruments		1.02e-07		0.00133		6.75e-08		0.00677

Notes: *Significant at 10% level **significant at 5% level *** significant at 1% level. P-values in square brackets

Table 6 shows the results for revenue per head, our measure of labour productivity, as the dependent variable. Once again the IV estimates of β are reasonably well-determined and much larger in absolute value than their OLS analogues although, in contrast to the capital-per-head regressions, the latter are not statistically significant. The coefficients do get somewhat smaller over time, perhaps as firms have more time to adjust their factors of production. The OLS estimates of β_{1t} are economically very small and statistically insignificant. In contrast, our IV estimates are economically large and relatively precisely estimated. Once again, diagnostic suggests any bias from weak or invalid instruments are unlikely to be large.

Interpreting the numbers, columns (2) and (6) shows that a credit shock that reduces a firm's level of borrowing by 10% by 2008 or 2009 reduced labour productivity by 8% and 5-6% respectively. These are of a similar size to the effect of credit on capital per worker. A standard neoclassical production function at the firm level would predict an effect on productivity equal the product of the effect on capital per worker and the capital share (around one-third). The size of the productivity effect is so large that impaired credit also reduced either total factor productivity or factor utilisation. Tables 7 and 8 suggest that most of the effect comes through reduced revenue, with little adjustment in employment.

Table 9 shows our estimate of the impact on average pay. As above, the IV estimates are much larger in absolute value than their OLS analogues, and suggest that a firm-level credit shock reduces wages in approximately the same proportion as it reduces productivity. This suggests that firms were able to share some of the costs from the credit shock with

Table 7 Second stage - baseline sample: Revenue. P-values in square brackets. The table shows results from the second stage regression of the change in revenue per head, Δy_{it} , on the change in debt, $\Delta debt_{it}$, which is instrumented using a vector of bank dummies in 2007. Changes in debt and y are calculated using symmetric growth rates. Each column represents a cross-sectional regression for the change in time t relative to 2007, where t is (2008, 2009). Regressions control for industry sector, firm entity type, age and log level of revenue in 2007. Balanced are for firms present in both years. Firms that have become inactive or are in liquidation are included in the sample as having zero revenue. Failed firms are included as having zero revenue and zero assets

	Unbalanced				Balanced			
	2008		2009		2008		2009	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in debt	0.054*** [0.01]	0.918*** [0.16]	0.086*** [0.01]	0.379** [0.16]	0.058*** [0.01]	0.906*** [0.16]	0.093*** [0.01]	0.411** [0.18]
Observations	12,959	13,444	11,106	11,418	10,155	10,471	10,191	10,471
R-squared	0.036	-1.697	0.061	-0.115	0.037	-1.802	0.062	-0.155
Sargan		6.909		7.066		6.766		8.719
Sargan pval		0.227		0.216		0.239		0.121
Kleinbergen-Paap		7.227		3.632		7.381		2.967
P-val of excl. instrum.		1.02e-07		0.00133		6.75e-08		0.00677

Notes: *Significant at 10% level **significant at 5% level *** significant at 1% level

their employees. It also suggests these firms were operating in a labour market that was non-Walrasian to an important extent, such that wages were not equalised across similar employees but differed ex post according to the credit supply experienced by their employers. Comparing them to Table 6, it seems that firms in our sample cut wages by roughly the same degree as labour productivity, supporting the notion that workers are paid their marginal product. In line with the results above, the diagnostic tests suggest that our instruments are strong and valid in 2008, but less so in 2009.

When comparing IV estimates across our three response variables, the estimated parameters are typically closer to zero in 2009 than in 2008, and our instruments are typically weaker and closer to being invalid. There are four possible reasons, two econometric and two economic, as to why this is the case. First, the instruments are weaker for 2009, which may be one reason why the estimated coefficient is smaller in absolute value (i.e. it may be biased towards the OLS parameter). Second, the sample shrinks between the two years, in part because some of the firms in the 2008 sample had gone bankrupt or become inactive by 2009. In the (likely) event that some of this attrition is non-random with respect to credit supply, the effect we observe among surviving firms will be different; it is noteworthy in this regard that the difference between the estimates for 2008 and 2009 is smaller on the balanced sample. Appendix D below shows that sample attrition may invalidate our Hansen-Sargan statistics, leading to spurious rejections of the null of valid instruments. Third, the effect of a persistent credit shock on productivity may fade over time, as firms have more time to reorganise production or finance expenditures through other means. Fourth, the dispersion of credit supply across banks may have fallen in those latter years.

Table 8 Second stage - baseline sample: Employment in heads. P-values in square brackets. The table shows results from the second stage regression of the change in revenue per head, Δy_{it} , on the change in debt, $\Delta debt_{it}$, which is instrumented using a vector of bank dummies in 2007. Changes in debt and y are calculated using symmetric growth rates. Each column represents a cross-sectional regression for the change in time t relative to 2007, where t is (2008, 2009). Regressions control for industry sector, firm entity type, age and log level of revenue in 2007. Balanced are for firms present in both years. Firms that have become inactive or are in liquidation are included in the sample as having zero revenue. Failed firms are included as having zero revenue and zero assets

	Unbalanced				Balanced			
	2008		2009		2008		2009	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in debt	0.053*** [0.01]	0.107 [0.07]	0.089*** [0.01]	-0.131 [0.13]	0.057*** [0.01]	0.115 [0.07]	0.088*** [0.01]	-0.171 [0.15]
Observations	12,959	13,444	11,106	11,418	10,155	10,471	10,191	10,471
R-squared	0.027	0.015	0.052	-0.109	0.028	0.015	0.051	-0.173
Sargan		4.137		5.333		3.942		5.546
Sargan pval		0.530		0.377		0.558		0.353
Kleinbergen-Paap		7.227		3.632		7.381		2.967
P-val of excl. instrum.		1.02e-07		0.00133		6.75e-08		0.00677

Notes: *Significant at 10% level **significant at 5% level *** significant at 1% level

The top 4 UK banks issued £120 billion in equity in 2008 - 2009, but not after that. As a result, the pressure on banks to cut back on lending, may have eased in 2010 and 2011 as their capital positions started to improve. Consistent with this, when we extend our analysis for 2010 and 2011, we find that our results completely fade: the instrument is measured as weak and the estimates of β_1 in the second stage are poorly determined and widely dispersed.

6.2 Firm survival

So far, the analysis has focussed on the intensive margin. However, an important question is whether the attrition in our sample due to firms ceasing to operate is random or not with respect to the other variables in our model. It seems likely that changes in credit supply will influence firm survival. Indeed, our dataset allows us to quantify the impact of credit supply shocks on firm survival.

The right hand side of our regression model only includes *time-invariant* variables observed in 2007. This means that we can evaluate the predicted change in a firm's borrowing among dead firms, in other words how much a firm with similar pre-crisis characteristics would have been expected to borrow had it survived. This is a natural metric with which to assess the impact of predicted borrowing, and in particular the contribution of bank identity to it, as a determinant of firm survival. We construct a binary cumulative failure indicator f_{it} taking the value of zero if firm i is alive in year t and 1 if the firm failed in or before

Table 9 Second stage - baseline sample: Average pay. The table shows results from the second stage regression of the change in total remuneration per head, Δy_{it} , on the change in debt, $\Delta debt_{it}$, which is instrumented using a vector of bank dummies in 2007. Changes in debt and y are calculated using symmetric growth rates. Each column represents a cross-sectional regression for the change in time t relative to 2007, where t is (2008, 2009). Regressions control for industry sector, firm entity type, age and log level of revenue in 2007. Balanced are for firms present in both years. Firms that have become inactive or are in liquidation are included in the sample as having zero revenue. Failed firms are included as having zero revenue and zero assets

	Unbalanced				Balanced			
	2008		2009		2008		2009	
	OLS	IV	OLS	IV	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in debt	0.004 [0.00]	0.851*** [0.15]	-0.007 [0.00]	0.670*** [0.18]	-0.000 [0.00]	0.858*** [0.15]	-0.002 [0.00]	0.764*** [0.22]
Controls	YES	YES	YES	YES	YES	YES	YES	YES
Observations	13,444	13,444	11,418	11,418	10,471	10,471	10,471	10,471
Hansen-Sargan Stat		3.305		12.55		5.496		11.73
Hansen-Sargan P-values		0.653		0.0279		0.358		0.0387
Kleinbergen-Paap		7.227		3.632		7.381		2.967
P-value of excl. instruments		1.02e-07		0.00133		6.75e-08		0.00677

Notes: *Significant at 10% level **significant at 5% level *** significant at 1% level. P-values in square brackets

year t .¹⁰ We then run logit regressions of firm failure or survival on the predicted value of credit supply, plus the non-bank controls in Eq. 2

$$f_{it}^* = \alpha_t \widehat{\Delta d}_i + \beta_t x_i + u_{it} \quad (6)$$

$$Pr(f_{it}^* > 0) = F(\alpha \widehat{\Delta d}_i + \beta x_i) \quad (7)$$

The estimated coefficient on credit supply will therefore capture the effect that credit supply, as identified with bank ID and measured in units of credit volumes, has on bank survival. A negative coefficient would mean that firms which would have been able to borrow more, had they survived, would have been less likely to fail.¹¹

Table 10 sets out our estimates of the parameter α in Eq. 7, the effect of a reduction in credit supply on the probability of firm death.¹² The results are negative and significant

¹⁰We create a proxy for firm death by looking at whether a firm's status (when the data were collected) is not 'active' and looking for the first year in which balance-sheet data such as assets are either zero or missing: we assume the firm failed in that year.

¹¹In principle, a reasonable alternative would be to estimate responses for both continuous left hand side variables and firm survival jointly, e.g. through a Heckman-type model. The problem is that with our data we lack an additional exclusion restriction and would need to rely on distributional assumptions and the curvature of the inverse Mills ratio for identification. For this reason, we focus here on the question of whether factors that predict credit supply among surviving firms also predict whether a firm survives at all.

¹²We calculate standard errors and significance levels for the firm survival model (6) and (7) by bootstrap resampling with replacement from our full dataset. This is necessary because the samples for Eqs. 6 and (7) are different, such that normal IV formulae for standard errors are not applicable.

Table 10 Firm survival logit regression - baseline sample. The table shows the estimated coefficients for α in Eq. 5

	(1) 2008	(2) 2009
Predicted debt	-6.26*** [1.08]	-5.27*** [0.82]
Controls	YES	YES
Observations	82,624	82,709

for all years under consideration. A negative parameter means that higher predicted credit supply (the amount a firm's borrowing would have changed, conditional on survival, based on its banking relationships and other observables) increases the probability that the firm survived up until the period in question. This suggests that a widespread contraction in credit supply will tend to increase corporate insolvencies across the economy, an effect we quantify in Section 7.

6.3 Robustness

6.3.1 Placebo tests

Our identifying strategy works by exploiting variation across banks in the changes in terms on which they supplied credit during the financial crisis. This variation was likely to have been relatively large during the period of funding and credit market turbulence, asset-price volatility and bank nationalisation of 2008-9. Conversely, if our instruments are valid and banking relationships are randomly assigned with respect to corporate outcomes, our instruments are more likely to be only weakly relevant at a time of tranquil market conditions. But if there is also endogenous variation in bank relationships that is relevant for credit supply and corporate performance, violating our identifying assumptions, then our instruments could turn out to be relevant when they should not be.

With this in mind, we re-run the regressions above but substitute 2005 in place of 2007 for our sample selection rule¹³ and as the base year against which changes in debt, labour productivity and so on are measured.

Table 11 shows the first stage regression results, analogous to Table 4. Table 12 presents the results from the coefficient on the $\Delta debt_t - 2005$, analogous to the first row of the balanced regressions seen in Tables 5-9. In contrast to our baseline results, the F-statistics for the first-stage regressions and the associated IV estimates of β are typically insignificant at standard levels. The only exception to this is when revenue per head is the dependent variable and the change is measured between 2005 and 2007, in which case the coefficient on change in debt is significant at the 10% level. This could be a false positive, or could suggest that the market turbulence experienced in the second half of 2007 may already have been having an effect on the real economy. However, in no case are our instruments relevant in 2006, a time of uniformly easy funding conditions across UK banks.

¹³We select firms in the Bureau Van Dijk database who had a positive level of debt in 2005 rather than 2007

Table 11 First stage - placebo regressions. The table shows results from regressions on the change in debt, $\Delta debt_{it}$, on a vector of bank dummies in 2005. $\Delta debt_{it}$ is calculated using symmetric growth rates. Each column represents a cross-sectional regression for the change in time t relative to 2005, where t is (2006, 2007). Regressions control for industry sector, firm entity type, age and log level of revenue in 2005. Balanced are for firms present in both years

	Unbalanced		Balanced	
	2006 (1)	2007 (2)	2006 (3)	2007 (4)
Bank 1	-0.021 [0.02]	-0.049** [0.02]	-0.020 [0.02]	-0.046** [0.02]
Bank 2	0.019 [0.02]	-0.024 [0.03]	0.010 [0.02]	-0.030 [0.03]
Bank 3	-0.001 [0.02]	-0.023 [0.02]	0.006 [0.02]	-0.021 [0.02]
Bank 4	-0.013 [0.01]	-0.030 [0.02]	0.000 [0.02]	-0.026 [0.02]
Other financial institution	0.001 [0.03]	-0.098*** [0.04]	-0.004 [0.03]	-0.103*** [0.04]
Non financial institution	-0.015 [0.02]	-0.025 [0.02]	-0.008 [0.02]	-0.025 [0.02]
Controls	YES	YES	YES	YES
Observations	13,862	11,266	10,532	10,532
F statistic	1.152	2.392	0.914	2.451
P-value of regression	0.269	8.74e-05	0.590	5.39e-05
P-value of all banks	0.655	0.127	0.849	0.139

Notes: *Significant at 10% level **significant at 5% level *** significant at 1% level. P-values in square brackets

6.3.2 Difference-in-differences

To test the robustness of our findings, in this section we employ an alternative difference-in-differences approach.

A reasonable alternative to our baseline specification would be to group banks according to whether they are 'strong' or 'weak', and therefore more or less likely to provide credit, and then perform a difference-in-differences analysis comparing firms who have relationships with these two groups.

The challenge with this approach is that it is not obvious how to group banks. On one hand, banks like RBS and Lloyds became so weak that they were nationalised. On the other,

Table 12 Second stage - placebo regressions. The table shows results from the second stage regressions of the change in Δy_{it} , on the change in debt, $\Delta debt_{it}$, which is instrumented using a vector of bank dummies in 2005. Changes in debt and y are calculated using symmetric growth rates. Each column represents a cross-sectional regression for the change in time t relative to 2005, where t is (2006, 2007). Regressions control for industry sector, firm entity type, age and log level of revenue in 2005. Balanced are for firms present in both years. Firms that have become inactive or are in liquidation are included in the sample as having zero revenue. Failed firms are included as having zero revenue and zero assets

Δy_{t-2005} :	2006		2007	
	OLS (1)	IV (2)	OLS (3)	IV (4)
Capital per head	0.137*** [0.01]	0.503 [0.44]	0.148*** [0.01]	-0.161 [0.31]
Revenue per head	-0.011 [0.01]	0.239 [0.34]	0.004 [0.01]	0.525* [0.28]
Average pay	-0.006 [0.01]	0.290 [0.31]	-0.008* [0.00]	0.100 [0.20]
Controls	YES	YES	YES	YES
Observations	10,421	10,421	10,421	10,421
P-value of instruments		0.791		0.163
Kleinbergen-Paap		0.954		2.492

Notes: *Significant at 10% level **significant at 5% level *** significant at 1% level. P-values in square brackets

nationalisation itself may have prompted a change in lending policy and actually boosted credit supply from the affected banks (see Rose and Wieladek (2014) for evidence that nationalisation affected the lending of UK banks).

Furthermore, banks facing tougher funding conditions may have had more of an incentive to provide loan support or forbearance to weaker firms to avoid having to realise further losses on their balance sheet. This might have had a positive impact on measured lending. For example, latest estimates suggest the scale of forbearance across non-CRE SME borrowers is likely to have been relatively small by 2013, but could have been higher immediately after the crisis Arrowsmith et al. (2013).

To overcome this issue we compare the average cumulative growth rate of total debt since the crisis, controlling for other observable differences, across firms that borrowed from each bank. Those with an average growth rate above the median are defined as strong and the rest as weak. We then construct a dummy variable for whether a bank is 'weak' or not, and run OLS regressions of the following form:

$$\Delta y_{it} = \beta_{0t} + \beta_{1t}D_{it} + \beta_{2t}x_{it} + \epsilon_{it} \quad (8)$$

where i indexes firms, t is time, D is a dummy variable equal to 1 for 'weak' banks, y is the response variable of interest and x is a vector of observable firm characteristics.




Fig. 5 Differences in debt and productivity growth across firms. The figures show the coefficient on a dummy variable (from Eq. 8) indicating relationships with a strong or weak bank in a regression of firm debt or productivity on banking relationships and other controls

Figure 5 outlines the results and plots the coefficient estimates for the dummy variable for each year, as well as the 95% confidence intervals around it. Panel (a) shows that the growth rate of debt among firms with weak banks was significantly lower than among those with strong banks. Panel (b) shows the difference in productivity growth between the two groups of firms. Productivity grew significantly less quickly among those firms whose banks cut lending more.

7 Economic magnitude

This section provides a set of simple illustrations to demonstrate the economic significance of our results, by translating our microeconomic findings into macroeconomic numbers. These calculations, although only illustrative, imply that the impact of the credit shock on the macroeconomy may have been substantial.

If you combine our results with an estimate of how much of the fall in firms' borrowing was due to changes in credit supply, it is possible to explain a large slice of the fall in investment and productivity.




Fig. 6 Estimate of the impact of credit supply on labour productivity. The pre-crisis trend is calculated between 1997Q1 and 2007Q2. The grey swathe assumes that 50% of the deviation in PNFC bank credit relative to trend was due to changes in credit supply. The upper bound for the impact on productivity uses an elasticity of 0.8 and the lower bound 0.5, based on the estimates in Table 6. Source: ONS, Author calculations

Suppose, for the sake of illustration, that half of the fall in the amount firms borrowed was due to less credit supplied by banks, and half because they demanded less from banks. Indeed, recent research by the Bank of England suggests that the vast majority of the initial fall in aggregate lending was due to credit supply shocks, rather than decreased demand (Barnett and Thomas (2014)). We multiply this by the proportion of firms with debt, which is 60% according to the BvD FAME database in 2007. We then multiply this by the relevant elasticity estimates from our regressions to derive some simple predictions of the aggregate impact of the credit shock.

Figure 6 presents the results for productivity, and Table 13 shows the results for productivity, capital per head and wages.

These estimates suggest that the credit supply shock could have lowered the level of labour productivity by 5–8% by 2013. This compares to an aggregate shortfall of 17% by 2013. Barnett et al. (2014a) argue that by the end of 2013, structural factors - including reduced investment in physical and intangible capital, and impaired resource allocation - could explain around six to nine percentage points of the deviation relative to pre-crisis trend. Our results are in a similar ballpark, and suggest that the credit supply shock could have been a major underlying cause behind these factors.

Table 13 Economic significance - baseline regressions

Calculation step	Capital per head %	Revenue per head %	Average pay %
Corporate debt relative to pre-crisis trend, 2013 (a)	−34	−34	−34
Proportion of fall in debt due to credit supply, 2013 (b)	50	50	50
Impact of 1% change in credit supply (c)	0.5 to 0.6	0.5 to 0.8	0.7 to 0.9
Proportion of firms with bank debt (d)	60	60	60
Predicted aggregate impact, 2013 (a)x(b)x(c)x(d)	−5 to −6	−5 to −8	−7 to −9
Actual deviation relative to pre-crisis trend, 2013	−6	−17	−27

Source: ONS; authors calculations

Table 14 Economic significance - firm survival

Calculation step	Percent
Change in corporate debt, 2007 to 2009 (a)	-8
Impact of 1% change in credit supply on probability of bankruptcy (b)	6
Proportion of firms with bank debt (c)	60
Predicted pickup in liquidations, (a)x(b)x(c)	27
Actual pickup in liquidations, 2007 to 2009	59

Source: ONS; authors calculations

These estimates predict that the level of capital per worker would be 5-6% lower by 2013. One serious consideration to bear in mind here is comparability of the data. The measure of capital used in our data is the level of tangible assets among non-CRE PNFCs. The aggregate capital stock includes a wider variety of assets across all PNFCs. Nonetheless, they suggest that by 2013 a large part of the 6% deviation in capital per worker may have been driven by reduced credit supply.

These back-of-the-envelope calculations also suggest that the credit shock would have contributed to lower levels of average pay - by around 7-9% in 2013 or up to a third of the deviation in real wages relative to their pre-crisis trend.

These estimates rely heavily on the assumption that the effects of credit supply are persistent over time. As discussed earlier, we only find significant effects for 2008 and 2009 and cannot distinguish whether this is because our instrument becomes less relevant over time (firms eventually switch banks) or whether the impact of the credit supply shock fades.

Turning to firm survival, in a logistic regression, such as that defined by Eq. 7, the slope parameter α implies that a one-unit increase in the explanatory variable increases the odds of the event by a factor e^α . So our estimate of the parameter α in Eq. 7 of around -5 suggests that a credit supply shock that would have reduced borrowing by 10% would raise the odds of bankruptcy by about $e^{-5 \times 10\%} \approx 60\%$.

Corporate liquidations in England and Wales rose by around 60% between 2007 and 2009, while the stock of corporate debt fell by 7.5%. Assuming again that the fall in corporate debt over this period represented a shock to credit supply, then multiplying these numbers with the proportion of firms with bank debt suggests that the credit shock may have accounted for almost half of the pick-up in company liquidations. This calculation is illustrated in Table 14.

The estimates provided above are purely illustrative, yet they suggest that the size of our parameter estimates in Section 5 are not only statistically significant but also economically large.

8 Conclusion

This paper examines the relationship between credit supply and corporate outcomes in the UK over the financial crisis.

We find that firms facing a reduction in credit supply experienced greater falls in labour productivity and capital per worker. We argue that this is due to an increase in the shadow price of capital causing firms to substitute towards more labour-intensive technologies in production. Our results suggest that a 10% contraction in credit supply led, on average, to a

5–6% fall in capital per worker and a 5–8% in labour productivity. The estimated impact on labour productivity is large, suggesting lower capital investment may have been associated with lower levels of innovation and technological development.

Our results also suggest average pay fell further in firms more exposed to the credit shock, and in similar proportion to labour productivity, even though these firms were hiring labour in the same markets as less exposed firms. We find that a 10% fall in credit supply led, on average, to a 7–9% fall in average pay for the firms affected. This suggests firms were able to share the impact of the credit shock with workers, through lower wage growth.

We find that firms facing adverse credit supply shocks were more likely to fail. Our results predict that a 10% decrease in credit supply would increase the probability of bankruptcy by 60%.

These parameter estimates are both statistically significant and economically large. They suggest that the credit supply shock caused by the recent financial crisis may explain around 5–8 percentage points of the 17% shortfall in labour productivity relative to its pre-crisis trend by 2013, half of the shortfall in wages, and nearly half of the increase in company liquidations between 2007 and 2009.

Our identification strategy relies on pre-crisis banking relationships that decay non-randomly over time. A key limitation of this empirical design is that we are unable to say how persistent these effects might be. Although we are only able to identify the impact of the contraction in credit supply in 2008 and 2009, it may be that either the shock was more persistent or the effects on the corporate sector longer lasting. For example, it could be that firms are somehow permanently scarred by temporary credit shocks, or it could be that they are able to catch up to the counterfactual pre-crisis trend path once the credit shock abates. Recent work by Mariano and Gine (2015) finds that the outcomes for firms in distress depends on the actions of its creditors. They look at what happens when a debt covenant is violated, a situation where the creditors acquire some control rights over the firm. They find that future investment declines when the firm has few growth opportunities but it may actually increase otherwise, as creditors consider the benefits of growth opportunities to meet outstanding debt obligations.

Our estimates are far from perfect. Indeed the variation in lending post-crisis between banks in our first stage regression is not as large as you might expect. However, our instruments pass a set of standard tests and are consistent with the results from studies based on other countries.

This study does not consider what is happening on the extensive margin. It only tracks a cohort of firms who had debt in 2007, and therefore does not examine the outcomes of firms born after 2007 or firms applying for new borrowing after 2007. This means we cannot say anything about how firm entry is affected by credit supply, or what happens to the factors that become unemployed when firms fail, either during or after the period of crisis. Barnett et al. (2014c) suggests impaired capital allocation between firms may have been an important driver of the weakness in productivity. There are also likely to have been factors, not covered in this study, helping to keep unproductive firms alive after the crisis, including forbearance by banks and the tax authorities, and low levels of interest rates Barnett et al. (2014b). Overall, the durability of the productivity slowdown since the crisis make this an important avenue for future research.

Acknowledgments The authors would like to thank Orazio Attanasio, Richard Button, Matthieu Chavaz, Kieran Dent, Antonio Guarino, Neeltje van Horen, Anil Kashyap, Veronica Rappoport, Isabelle Roland, Andrew Rose, Silvana Tenreyro, Ryland Thomas, Alex Tuckett, John van Reenen, Arzu Uluc, Garry Young, an anonymous referee, and seminar participants at the Bank of England, LBS, LSE, UCL, Oxford, University

of York, IAAE 2015, Institute for Fiscal Studies conference and CEPR-CFM conference for helpful comments. We would like to thank James Barker for excellent research assistance. Any views expressed are solely those of the authors and so cannot be taken to represent those of the Bank of England or members of its Committees or to state Bank of England policy. All errors are our own.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix A: The relationship between banks funding costs and exposure to CRE firms

Commercial property played an important role during the financial crisis. The rapid pickup in debt tied to commercial property investments prior to the crisis contributed to the rapid acceleration in prices, and the subsequent falls in prices led to a sharp rise in non-performing loans. While debt write-offs on CRE loans picked up to around 2% after the crisis, this is likely to significantly understate the scale of non-performing CRE loans on bank's balance sheets. Recent Bank of England work suggests the median rate of forbearance on CRE loans across banks reached 20%, with the worst performing bank portfolio reaching nearly 50% (Benford and Burrows 2013).

Figure 7 compares the share of UK CRE loans for big five UK banks in 2007 with the change in funding costs, proxied by the change in CDS premia between 2007 and 2009. The positive correlation supports the notion that a bank's exposure to CRE was correlated with the subsequent pickup in its funding costs. For this reason, and in line with Bentolila et al. (2013), we exclude CRE firms from our sample.

Figure 8 compares the change in funding costs with each banks exposure to the non-CRE corporate sector. The lack of an obvious relationship is consistent with the premise




Fig. 7 Bank exposure to CRE and change in CDS premia. The proportion of commercial real estate (CRE) lending from each bank is calculated as the stock of loans and advances to the UK non-residential property sector relative to the total loan book (covering all countries). The non-residential property sector here includes lending to the construction sector - to ensure consistency in reporting across banks. Data for HSBC are for for total European CRE lending as a proportion of the total loan book (published UK data were not available). Source: Published annual accounts for Barclays, HSBC, Lloyds, HBOS and RBS in 2007; MarkIT; authors calculations




Fig. 8 Bank exposure to non-CRE PNFCs and change in CDS premia. The proportion of private non-financial corporation (PNFC) lending from each bank is calculated as the stock of loans and advances to all industry sectors (excluding finance commercial real estate (CRE)) relative to the total loan book (covering all countries). Data for HSBC are for total European PNFC lending as a proportion of the total loan book (published UK data were not available). Data for Lloyds include loans to financial services (which were grouped together with business services in their annual published accounts). Source: Published annual accounts for Barclays, HSBC, Lloyds, HBOS and RBS in 2007; MarkIT; authors calculations

that banks non-CRE corporate loan book did not cause the variation in funding costs and balance sheet health in the immediate aftermath of the crisis.

Appendix B: OLS Bias

To see the determinants of the OLS bias, consider a very simple model of borrowing, d , output, y , and credit supply (perfectly elastic at interest rate r) at the firm level.

$$d = \alpha_1 y + \alpha_2 r + \epsilon_d \tag{Credit demand}$$

$$r = \epsilon_s \tag{Credit supply}$$

$$y = \gamma_1 \epsilon_s + \epsilon_y \tag{Output}$$

In terms of the exogenous shocks of the model, realised output and borrowing are given by

$$\begin{aligned} d &= \alpha_1 (\gamma_1 \epsilon_s + \epsilon_y) + \alpha_2 \epsilon_s + \epsilon_d \\ &= (\alpha_1 \gamma_1 + \alpha_2) \epsilon_s + \alpha_1 \epsilon_y + \epsilon_d \\ y &= \gamma_1 \epsilon_s + \epsilon_y \end{aligned}$$

If we were simply to regress output on credit volumes, the expected value of our OLS parameter estimate would be

$$\begin{aligned} E \left[\hat{\beta}_1^{OLS} \right] &= \frac{cov(y, d)}{var(d)} \\ &= \gamma_1 \frac{\sigma_s^2 (\alpha_1 \gamma_1 + \alpha_2)}{(\alpha_1 \gamma_1 + \alpha_2)^2 \sigma_s^2 + \alpha_1^2 \sigma_y^2 + \sigma_d^2} + \frac{\alpha_1 \sigma_y^2}{(\alpha_1 \gamma_1 + \alpha_2)^2 \sigma_s^2 + \alpha_1^2 \sigma_y^2 + \sigma_d^2} \end{aligned}$$

The OLS estimator is biased for two reasons. First, the bias is an increasing function of $\alpha_1 \sigma_y^2$ - i.e. output shocks will bias the OLS parameter upwards to the extent that borrowing is an increasing function of output and that there are output shocks in the sample. Intuitively,

if credit volumes are strongly increasing in output, and output varies autonomously a great deal, the parameter estimate in an OLS regression of output on credit volumes will be biased upwards. Secondly, credit demand shocks will bias the parameter towards zero, as they will raise the variance of the right-hand side variable in the regression.

For this reason, we adopt an instrumental variables approach. For each of the firms in our sample, we have information about the banking relationship between them and different banks. The banking relationship b is correlated with credit supply but uncorrelated with any of the other shocks, such that $b = \mu\epsilon_s + \epsilon_b$. In expectation, our IV estimator is then

$$\begin{aligned} E[\hat{\beta}_1^{IV}] &= \frac{cov(y, b)}{cov(d, b)} \\ &= \frac{E[(\gamma_1\epsilon_s + \epsilon_y)(\mu\epsilon_s + \epsilon_b)]}{E[(\alpha_1\gamma_1 + \alpha_2)\epsilon_s + \alpha_1\epsilon_y + \epsilon_d](\mu\epsilon_s + \epsilon_b)} \\ &= \frac{\gamma_1\mu\sigma_s^2}{\mu(\alpha_1\gamma_1 + \alpha_2)\sigma_s^2} \\ &= \frac{\gamma_1}{(\alpha_1\gamma_1 + \alpha_2)} \end{aligned}$$

This coefficient provides an unbiased estimate of the effect of credit supply shocks on output normalised by their effect on borrowing.

Appendix C: Attrition and overidentification tests

To the extent that credit supply affects firm survival, it may also affect the results of the standard tests of overidentification restrictions used to assess instrument validity. Standard tests of overidentifying restrictions work by looking how far from zero our sample analogues $b'\hat{\epsilon}$ of the population moment conditions $E[\epsilon_i b_i] = 0$ are. However, if there is nonrandom attrition in our sample, tests based on such overidentifying restrictions are unlikely to work. In particular, if the disturbances in the observation equation ϵ are correlated with those in the selection equation u , and our instrument d is a determinant of the latent selection variable f^* (this is a necessary condition for identification), then in general our instrument will be correlated with the disturbances in the observation equation, conditional on them being observed, i.e.

$$E[b'\epsilon | f^* > 0] \neq 0$$

To see this, assume that the moment condition $E[b'\epsilon] = 0$ holds in the population. However, what we observe is instead the sample analogue of $E[b'\epsilon | f^* > 0]$

$$\begin{aligned} E[b'\epsilon | f^* > 0] &= E_b[E[b'\epsilon | f^* > 0, b]] \\ &= E_b[b'E[\epsilon | f^* > 0, b]] \\ &= E_b[b'E[\epsilon | \alpha_t \widehat{\Delta d}_i + \beta_t x_i + u_{it} > 0, b]] \\ &= E_b[b'E[\epsilon | u_{it} > -\alpha_t \widehat{\Delta d}_i - \beta_t x_i, b]] \end{aligned}$$

Suppose for the sake of illustration that u and ϵ are both mean zero, and jointly normally distributed with covariance $\sigma_{u\epsilon}$. Then

$$E_b [b' E [\epsilon | u_{it} > -\alpha_t \widehat{\Delta d}_i - \beta_t x_i, b]] = E_b \left[b' \frac{\sigma_{u\epsilon}}{\sigma_\epsilon} \frac{\phi \left(\frac{-\alpha_t \widehat{\Delta d}_i - \beta_t x_i}{\sigma_\epsilon} \right)}{1 - \Phi \left(\frac{-\alpha_t \widehat{\Delta d}_i - \beta_t x_i}{\sigma_\epsilon} \right)} \right] \neq 0$$

So in general, even if our identifying restrictions hold, our instruments will be correlated with the residuals in the second-stage equations for the continuous variables. The intuition is that our instruments determine selection, so they will be systematically related to the unobserved variables u in the selection equation *among surviving firms*. The latter will be related systematically to the unobserved variables ϵ to the extent that firms which are more likely to survive are also more likely to invest, hire labour, and so forth. This means that our instruments can be correlated with the residuals of the second-stage equation, invalidating tests of validity based on overidentifying restrictions. We may therefore reject the null of validity more frequently than indicated by the significance level of the test.

Appendix D: Variable definitions

- Total debt: The total amount of overdrafts, short-term loans and long-term debt in a given year.
- Capital and capital per worker: The total amount of tangible assets, and the total amount per employee.
- Revenue per head: Total revenue divided by the number of employees. This is our preferred measure of labour productivity, discussed in more detail below.
- Average pay: Total remuneration divided by the number of employees. This is our measure of average wages.
- Employment: Total number of employees in a given year.

References

- Amiti M, Weinstein D (2013) How much do bank shocks affect investment? Evidence from matched bank-firm loan data. Federal Reserve Bank of New York Staff Reports 604
- Arrowsmith M, Franklin J, Gregory D, Griffiths M, Wohlmann E, Young G (2013) SME forbearance and its implications for monetary and financial stability. Quarterly Bulletin, Bank of England
- Bank of England (2013) Trends in lending. Technical report, Bank of England
- Barnett A, Batten S, Chiu A, Franklin J, Sebastia-Barriel M (2014a) The UK productivity puzzle. Quarterly Bulletin, Bank of England
- Barnett A, Batten S, Sebastia-Barriel M, Chiu A, Franklin J (2014b) The productivity puzzle: a firm-level investigation into employment behaviour and resource allocation over the crisis. Working paper, Bank of England
- Barnett A, Broadbent B, Franklin J, Miller H (2014c) Impaired capital reallocation and productivity. National Institute Economic Review
- Barnett A, Thomas R (2014) Has weak lending and activity in the united kingdom been driven by credit supply shocks? Working paper, Bank of England
- Benford J, Burrows O (2013) Commercial property and financial stability. Quarterly Bulletin, Bank of England
- Bentolila S, Jansen M, Jiménez G, Ruano S (2013) When credit dries up: Job losses in the great recession. Discussion Paper 9776, CEPR

- British Business Bank (2014) Small business finance markets. Technical report, British Business Bank
- Cameron AC, Miller DL (2015) A practitioner's guide to cluster-robust inference. *J Hum Resour* 50(2): 317–372
- Chodorow-Reich G (2014) The employment effects of credit market disruptions: Firm-level evidence from the 2008–09 financial crisis. *Q J Econ* 129(1):1–59
- Edgerton J (2012). Credit supply and business investment during the great recession: Evidence from public records of equipment financing. mimeo
- Flannery M, Giacomini E, Wang X (2013) The effect of bank shocks on corporate financing and investment: Evidence from 2007–2009 financial crisis. mimeo
- FSA (2011) The failure of the Royal Bank of Scotland: Financial services authority board report
- Gilchrist S, Zakrajsek E (2012) Credit spreads and business cycle fluctuations. *Am Econ Rev* 102(4):1692–1720
- Greenstone M, Mas A (2012) Do credit market shocks affect the real economy? Quasi-experimental evidence from the great recession and 'normal' economic times. Technical Report 12-27, MIT Department of Economics Working Paper
- Kashyap AK, Stein JC (2000) What do a million observations on banks say about the transmission of monetary policy? *Am Econ Rev* 3:407–428
- Khwaja AI, Mian A (2008) Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *Am Econ Rev* 98:1413–1442
- Mariano B, Gine JAT (2015) Creditor intervention, investment, and growth opportunities. *J Financ Serv Res* 47(2):203–228
- Ongena S, Peydro J, Horen Nv (2013) Shocks abroad, pain at home? Bank-firm level evidence on the international transmission of financial shocks Discussion Paper 2013-040. Tilburg University, Center for Economic Research
- Paravisini D, Rappoport V, Schnabl P, Wolfenzon D (2015) Dissecting the effect of credit supply on trade: Evidence from matched credit-export data. *Rev Econ Stud* 82(1):333–359
- Pattani A, Vera G, Wackett J (2011) Going public: UK companies' use of capital markets. Bank of England Quarterly Bulletin
- PCBS (2013) An accident waiting to happen: the failure of HBOS
- Peek J, Rosengren ES (1997) The international transmission of financial shocks: the case of Japan. *Am Econ Rev* 87:495–505
- Peek J, Rosengren ES (2000) Collateral damage: Effects of the Japanese bank crisis on real activity in the United States. *Am Econ Rev* 90:30–45
- Rose AK, Wieladek T (2014) Financial protectionism? first evidence. *J Financ* 69(5):2127–2149
- Stock JH, Yogo M (2002) Testing for weak instruments in linear IV regression
- van der Schans D (2012) SME access to external finance. BIS Economics Paper 16, Department for Business Innovation & Skills
- Van Reenen J (1996) The creation and capture of rents: Wages and innovation in a panel of UK companies. *Q J Econ* 111(1):195–226

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.