BANK OF ENGLAND

Staff Working Paper No. 773

Mortgages, cash-flow shocks and local employment

Fergus Cumming

December 2018

BANK OF ENGLAND

Staff Working Paper No. 773
Mortgages, cash-flow shocks and local employment

Fergus Cumming[(1)]

Abstract

This paper quantifies the local impact of monetary policy through the cash-flow channel during the Crisis by combining novel micro datasets with near-universal coverage of UK mortgages and employment. I estimate that a reduction in mortgage payments equivalent to 1% of household income led to around a 5 percentage point increase in employment growth in non-tradable businesses the following year. But the spatial distribution of mortgage and labour market structures resulted in significant heterogeneity of this effect across the country. Taken at face value, the estimates suggest that the overall effect of accommodative monetary policy on total employment growth in 2010 varied by around 1.5 percentage points across regions.

Key words: Mortgages, interest rates, monetary policy, employment.

JEL classification: E21, E52, G21.

---

(1) Bank of England and Nuffield College, University of Oxford. Email: fergus@bankofengland.co.uk

This paper contains the views of the author and not necessarily those of the Bank of England or its committees. The Secure Research Service agrees that the figures and descriptions of results in this paper may be published but the ONS does not necessarily endorse any interpretation of their statistical data. All statistical results remain Crown Copyright. This work uses research datasets which may not exactly reproduce National Statistics aggregates. The author would like to thank, but not implicate, Saleem Bahaj, John Barrdear, Rhys Bidder, Ben Broadbent, James Cloyne, Andrew Dilnot, Martin Ellison, Angus Foulis, Andreas Fuster, Andy Haldane, Michael McMahon, David Miles, John Muellbauer, Ricardo Reis, Fergal Shortall, Lars Svensson, Johannes Stroebel, Paulo Surico and Silvana Tenreyro. Thanks also to participants at the 2018 GSE Barcelona, 2018 IAAE and 2018 EEA conferences, and seminar participants at De Nederlandsche Bank, the Federal Reserve Board of Governors, OFCE and Oxford for helpful comments. All errors remain my own.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 3461 4030 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

Monetary policy works through a number of channels but little is understood about how the transmission mechanism varies across regions. One channel that has recently attracted attention is the shock to household disposable income that results from the interaction of policy rates and large debt contracts. This cash-flow effect has been considered by policymakers for some time (e.g., Bank of England (1999) and Bernanke (2007)) but there has been little attempt to test how it might amplify the overall responsiveness of households to monetary policy. I fill this gap by combining multiple datasets with near-universal coverage to investigate the importance of the cash-flow effect on aggregate employment and its spatial distribution.

I quantify the cash-flow effect on employment by exploiting cross-sectional variation in UK mortgage contracts and the dramatic fall in interest rates during the Great Recession. The 400 basis point reduction in policy interest rates between October 2008 and March 2009 initially benefitted some households more than others. Those holding large debt contracts with interest rates closely linked to benchmark rates received substantial favourable cash-flow shocks, which slackened liquidity constraints and boosted consumption.[1]

The central analysis is split into two stages. I first model the evolution of almost six million mortgage payment flows to estimate the household-level change in cash flows that followed from the systematic easing of monetary policy in the autumn of 2008. I go on to define locally non-tradable firms as those generating revenues from nearby customers, in the spirit of Mian and Sufi (2014) and Giroud and Mueller (2017), and group my households into neighborhoods. I then exploit the staggered timing of neighborhood-level interest rate pass-through to estimate the relationship between changes in mortgagor repayments and the growth of locally non-tradable employment during 2009, which proxies for changes in spending when granular consumption data are unavailable.

Two features of the UK mortgage market aid the quest for identification of this cash-flow effect. First, many people were as-good-as randomly assigned a fixed or variable-rate mortgage before October 2008. In the preceding decade the average length of a mortgage interest rate fix was around two years, which is very short by international standards. Short fixes and an active refinancing market before 2008 meant the decision to take a fixed or variable-rate mortgage was perceived to have a small impact on the lifetime cost of the loan. I show that observable characteristics between fixed and variable-rate mortgagors were similar and mortgage choice was often driven by whichever contract had the lowest initial interest rate, even when the interest rate wedge between fixed and variable-rate contracts reflected information embedded in the yield curve. By the same token, households did not base mortgage choices on an anticipation of future

[1] Even those that benefitted less from the cash-flow effect benefitted from looser monetary policy through other channels.

monetary policy action. Survey evidence shows that only 10% of households in August 2008 expected policy rates to fall substantially in the coming months. In sum, households chose their mortgages based on a number of factors that varied over time; almost 40% of remortgagors took out both fixed and variable-rate contracts in the immediate run-up to the Global Financial Crisis (henceforth, Crisis).

Second, the turbulence in the UK mortgage market, which accelerated after the failure of Lehman Brothers, restricted fixed-rate mortgagors' ability to refinance their contracts in order to benefit from lower interest rates. A combination of high early-repayment fees, lower collateral values and short fixation periods meant that most fixed-rate mortgagors waited for their interest rate to reset rather than actively seek out a new contract.[2] I estimate that in the UK at most 7% of people on fixed-rate contracts actively refinanced their mortgages in 2009 based on total remortgaging activity during this period. This stands in stark contrast to the US, where remortgaging spiked up following the monetary easing (e.g. see Beraja, Fuster, Hurst, and Vavra (2017)).

I document that the average household on a variable-rate mortgage contract experienced a reduction in repayments equivalent to 5% of gross income during the year after policy interest rates reached their 2009 floor.[3] Had interest rates instead increased, the cash-flow shock of course would have been negative. At the same time, around a quarter of mortgagors initially experienced no change in cash flows because their repayments were not directly tied to policy rates in the period after the monetary easing. I find that locally non-tradable establishments surrounded by a large fraction of variable-rate mortgagors increased their employment relative to areas awash with fixed-rate mortgagors. The effect disappears entirely when a placebo specification is run on establishments in the locally tradable sector, which are unlikely to generate a sizable fraction of their turnover from local residents.

Instead of relying on relatively large, contiguous administrative regions, I estimate average cash-flow shocks in the neighbourhoods immediately surrounding each of the 450,000 locally non-tradable establishments. People travel less than 10km for their average shopping trip so using standard administrative boundaries often misses key heterogeneity, especially when the proportion of fixed and variable-rate mortgages varies substantially within regions.[4] I therefore use a neighbourhood-aggregated version of the interest rate exposure channel developed in Auclert (2017) to estimate the employment effect of monetary policy at various levels of aggregation.

The distribution of mortgagors on different types of contract was a key driver for the spatial

[2]UK mortgage interest rates typically increase with the loan-to-value ratio, which ties collateral values to the cost of refinancing. This was especially true after 2008.

[3]Early repayment charges are on the order of 1% of the outstanding loan balance per year left to run of the fixation period. For the typical mortgage with an LTI of around 3, that usually makes early refinancing unattractive when other remortgaging costs are included.

[4]See Department for Transport (2016) for the UK or Kerr, Frank, Sallis, Saelens, Glanz, and Chapman (2012) for the US.

distribution of cash-flow shocks. By aggregating cash-flow shocks to the neighbourhood level I show that, using cross-sectional variation, a reduction in mortgage payments equivalent to 1% of household income led to around a 5 percentage point increase in employment growth in non-tradable businesses the following year. I use neighbourhood-aggregated micro data to control for an extensive set of neighbourhood characteristics including average income, age and leverage going into the Crisis. I also control for neighbourhood-specific changes in housing net worth and GVA between the summer of 2008 and the end of 2009 to rule out alternative channels of propagation. Over half of the overall effect comes from firms boosting staff numbers and the rest from establishments remaining open when they otherwise would have shut down.

A consequence of the cash-flow effect I find is that the joint spatial distribution of mortgage and labour market structures led to significant heterogeneity in the traction of conventional monetary policy across the country in the Great Recession.[5] Regions where there were a large number of variable-rate mortgagors were particularly sensitive to the direct effects of monetary policy in 2009 through this particular channel, and even more so when income gearing was high. This sensitivity was compounded in areas where the local economy was relatively self-contained and employed a large share of people in the locally non-tradable sector. I can use my household and establishment-level employment data to estimate the impact of the cash-flow channel of monetary policy on local employment markets at any level of aggregation I choose. My results suggest that the impact from the change in policy rates on one-year total employment growth in the Great Recession varied by around 1.5pp across local administrative regions. This regional heterogeneity diminished over time as short fixation periods meant that the vast majority of outstanding mortgages were linked to policy rates after 2010. Since both employment and mortgage market structures are likely to vary over time, their joint evolution is important for an up-to-date understanding of the monetary transmission mechanism.

My work contributes to the growing literature that cements the cash-flow effect via the mortgage market as a key transmission channel of monetary policy. The notion that the structure of the mortgage market might affect the sensitivity of consumption is well established (e.g., Rubio (2011); Calza, Monacelli, and Stracca (2013) and Cloyne, Ferreira, and Surico (2016)).[6] But there has recently been renewed attention on the benefits of variable-rate contracts, made more important with the observation that changes in disposable income might have large macroeconomic effects (e.g., Piskorski and Seru (2018); Guren, Krishnamurthy, and McQuade (2018) and Violante, Kaplan, and Weidner (2014)). The challenge has been to quantify the cash-flow effect in terms of macroeconomic variables. There is compelling evidence that cash-flow commitments, i.e. the debt service ratio, are important for household propensity for delinquency (e.g., Fuster

[5] This paper does not consider the effects of Quantitative Easing or any other unconventional policies.

[6] This is consistent with survey evidence that suggests many households who received a cash-flow windfall from lower mortgage repayments in 2009 increased consumption (Hellebrandt, Pezzini, Saleheen, and Williams (2009)).

and Willen (2013)) and outright default (e.g., Aron and Muellbauer (2016) and Byrne, Kelly, and O'Toole (2017)) but identification in this area is often restricted to a limited part of the mortgage market. I exploit the relatively balanced distribution of fixed and variable-rate contracts in the UK to estimate the national and regional employment effect of monetary policy through the cash-flow channel.

Identification of the microeconomic effects of monetary policy is often hampered by the endogeneity of an area's characteristics and the causal effect of a change in interest rates. Recent attempts have been made to overcome this by linking individual spending data to household balance sheets (e.g., Cava, Hughson, and Kaplan (2016); Flodén, Kilström, Sigurdsson, and Vestman (2017) and Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru, and Yao (2017)), but these approaches do not lend themselves as neatly to the macroeconomic or regional consequences of policy. In the case of the US, that is because the adjustable-rate mortgage market is a relatively small share of household secured lending. I therefore combine the loan-level approach with the more aggregated analysis of Mian and Sufi (2014) and Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru, and Yao (2017) to examine the employment growth of every locally non-tradable establishment in the UK and the evolution of household cash-flows immediately surrounding them. In addition, I use some well-known results that demonstrate how households struggle to make rational choices in the face of important financial decisions (e.g., Fornero, Monticone, and Trucchi (2011); Agarwal and Mazumder (2013) and Agarwal, Ben-David, and Yao (2017)). The combination of tightly defined neighbourhoods, microdata-aggregated controls and my fixed-variable-spread instrument means I am more confidently able to use cross-sectional variation to infer the causal impact of monetary policy on employment. I go on to show that although the mortgage market is important for understanding the monetary transmission mechanism through to employment, so is the make-up of local labour markets.

This paper also contributes to the literature surrounding the redistributionary consequences of monetary policy (e.g., Bullard (2014); Selezneva, Schneider, and Doepke (2015); Ozkan, Mitman, Karahan, and Hedlund (2016); Coibion, Gorodnichenko, Kueng, and Silvia (2017) and Bunn, Pugh, and Yeates (2018)), though the literature on regional heterogeneity is relatively sparse (e.g., Carlino and DeFina (1997) and Beraja, Fuster, Hurst, and Vavra (2017)). Perhaps the closest complementary paper to this one is work by Luck and Zimmermann (2017), which investigates a similar question for unconventional policy. But the vast bulk of UK mortgage lending is not conducted by regional banks and using mortgage heterogeneity allows identification of a more precise transmission channel. The employment effect I identify operates through changes in consumption so my work complements the early theoretical work of Jackman and Sutton (1982), supported by empirical evidence in Aron, Duca, Muellbauer, Murata, and Murphy (2012), and more recently Auclert (2017).[7] In particular, my results suggest a feedback link between Auclert's

[7]See Calza, Monacelli, and Stracca (2013) and Corsetti, Duarte, and Mann (2018) for studies examining

earnings heterogeneity channel and the *interest rate exposure channel* via the cash-flow effect when locally non-tradable employment is prevalent. In a similar spirit to Beraja, Fuster, Hurst, and Vavra (2017), I am able to track which regions are most responsive to changes in interest rates in the short run and go on to capture the knock-on effect monetary policy has on locally non-tradable employment at different levels of granularity.

The rest of this paper is organised as follows. Section 2 describes the data sources used for the information on firms, mortgages and other controls. Section 3 sets out the empirical strategy and clarifies the identification approach that runs throughout the whole paper. Section 4 presents the main results and Section 5 contains some important robustness checks. Section 6 sets the results in context and concludes.

2 Data

2.1 Mortgage data

My analysis uses the universe of residential mortgages issued by UK lenders since April 2005, collected by the Financial Conduct Authority (FCA) and distributed in the Product Sales Database (PSD).[8] It contains a wealth of information on property, borrower and lender characteristics at the time of origination. Using these mortgage flows I construct an estimate for the stock of mortgages as of July 2008, well before the failure of Lehman Brothers and the internationally coordinated policy interventions in the autumn that year. Appendix 7.1 goes through the steps needed to transform the mortgage flows into the stock.[9]

The 5.8m mortgages in my estimated stock form a large and representative sample of the residential mortgage market in 2008.[10] The UK mortgage market is broadly split into products that have a fixed interest rate at origination and those that have an interest rate linked to the Bank of England policy rate (Bank rate).[11] Mortgage terms tended to be between 25 and 30 years but the periods governing the path of the interest rate (henceforth, contractual maturity) have

heterogeneous effects of monetary policy through the housing channel.

[8]See http://www.fca.org.uk/firms/systems-reporting/product-sales-data/ for published high-level data. The PSD includes regulated mortgage contracts only, and therefore excludes other regulated home finance products such as home purchase plans and home reversions, and unregulated products such as second charge lending and buy-to-let mortgages.

[9]From 2015 the FCA began publishing the true stock of mortgages as part of the PSD making the process outlined in Appendix 7.1 much more straightforward for future studies. Data quality are increasing over time.

[10]There are no consistent estimates of the true universe of mortgages at this time but contemporary estimates of the stock suggest my sample is likely to represent around 80-85% of the relevant residential mortgages. Some industry bodies suggest that the true number is closer to nine million, though this includes second-charge mortgages. Regulatory data for 2015 shows that there were around 6.7m first-charge residential mortgages and Figure 15 of Appendix 7.1 provides an illustration of which mortgages I estimate may be missing from my stock.

[11]Before 2009 there were a small number of more exotic products such as mortgages with caps and floors, or with an interest rate linked to an alternative interest rate index.

historically been relatively short in the UK, especially compared to markets such as the US. The two most popular mortgages between 2005 and 2008 were those that had two-year variable or two-year fixed interest rates. So although I refer to the latter as *fixed*, they are actually much closer to a *short-run hybrid* mortgage using North American nomenclature. Short contractual maturities meant that the split between mortgagors on fixed and variable rates was always relatively even.

Following the end of the contractual maturity, interest rates revert to the so-called *Standard Variable Rate* (SVR). This is a mortgagee-set interest rate that loosely follows the path of Bank rate. Before the Crisis, the spread between the SVR and remortgage interest rates was around 300-400bp (with little variation across lenders), meaning it was usually beneficial for mortgagors to refinance upon contractual maturity. The majority of mortgagors therefore refinanced every two to three years during the Great Moderation.

One apparent concern with my estimated stock might be that the mortgages missing from the sample were in some way different to the others, biasing the results. In fact, because this study uses cross-sectional variation, bias is only likely to arise if the missing mortgages are somehow unevenly distributed across the country, which seems unlikely. Moreover, the flow of long-term fixed-rate contracts has always been very low in the UK. This means almost all the missing mortgages will have been linked to Bank rate in 2008, which means we can be confident about the *relative* proportions of fixed and variable-rate mortgages in a particular area. The proportions of variable and fixed-rate mortgages in the summer of 2008 are shown in Figure 1.

Figure 1: Distribution of Mortgages in July 2008

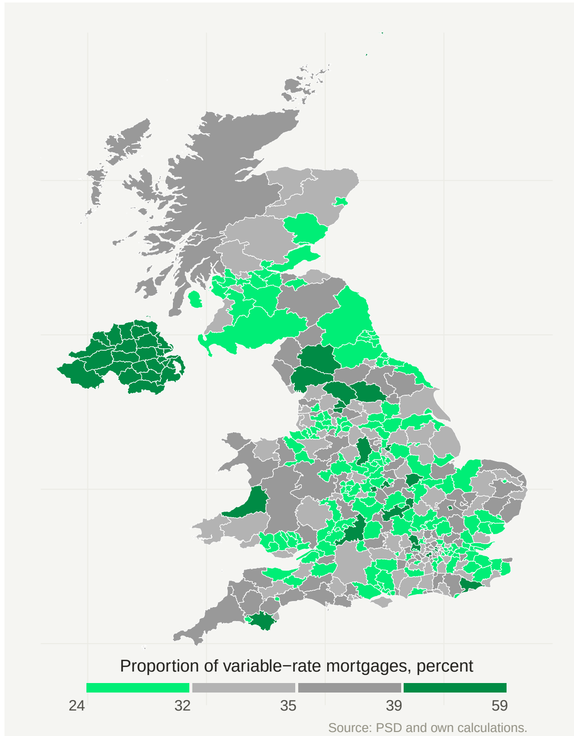

Figure shows proportion of variable-rate mortgages at origination. Colour breaks denote quartiles and 389 local authorities are shown.

I model how the associated characteristics evolved after origination. My primary interest is how monthly repayments responded to the monetary easing at the end of 2008. There are a wide range of mandatory fields in the PSD that have complete coverage of information collected at origination such as the date, location, borrower birth date, loan value, property value, household income and how the interest rate contractually varied over time. Other variables have less than complete coverage, often due to heterogeneous reporting practices of the mortgage lenders. Of these, the most important variable left blank is the the interest rate at origination. Fortunately, the highly competitive nature of the UK mortgage market means that I can accurately model what the likely interest rate would have been. In addition, the dramatic change in interest rates is more important than the level for the cross-sectional variation I use. Appendix 7.2 addresses how I deal with variables that have incomplete coverage in more detail.

I model mortgage cash-flows from the fourth quarter of 2008 through to the end of 2009. There are three possible types of borrower on the eve of the monetary easing in October 2008. Borrowers could either be part-way through a fixed-rate contract; part-way through a variable-rate contract; or beyond the initial contractual maturity (of either type). The first group experienced little or no cash-flow shock. The second received a substantial, favourable one. But the third group also often experienced a favourable cash-flow shock because the typical SVR fell in lock-step with mortgage rates at origination. Figure 2 shows the evolution of Bank rate, the two most common (new) mortgage contracts and the SVR.[12]

Figure 2: Mortgage Interest Rates

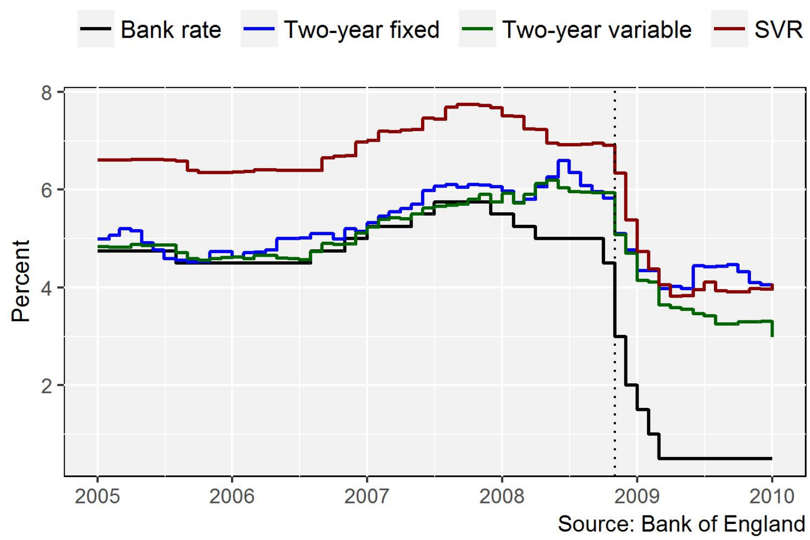

The first series shown is Bank rate. From the Bank of England quoted rates series (with identifying code in brackets): 2-year fixed 75% LTV mortgage (IUMBV34), 2-year variable 75% LTV mortgage (IUMBV48) and the Standard Variable Rate (IUMTLMV).

2.2 Employment data

The main source of employment data is the Business Structure Database (BSD), which contains financial information for over two million companies registered in the UK. The BSD is compiled as an annual snapshot from the Interdepartmental Business Register (IDBR), which requires firms to report information at the enterprise and local-unit level. Since the IDBR is based on Her Majesty's Revenue and Customs tax data, it captures the universe of economically active firms in the UK that are registered for income tax purposes. An enterprise is defined as the smallest combination of legal units that have a degree of autonomy from an enterprise group and can therefore be thought of as the overall business or firm. Enterprises are made up of one or more local units, or establishments, such as individual shops or restaurants. Businesses are required to report turnover at the enterprise level and employment at the establishment level, as well as geographic information for both.

To identify the effects of cash-flow shocks this paper follows Mian and Sufi (2014) and Giroud

[12] Deposit interest rates fell a similar amount at the end of 2008 but the overall offsetting cash-flow effect was small because household liabilities often dwarfed deposit assets. The second wave (2008-2010) of the UK Wealth and Assets Survey shows that the median mortgagor owed around £70,000 on their mortgage but only had around £1,000 in savings. Figure 2 shows that the SVR fell less than Bank-rate and so the cash-flow shock was larger for mortgagors on variable-rate contracts within their contractual maturity. This is taken account of in the cash-flow modelling.

and Mueller (2017) in categorising firms into tradable, *locally* non-tradable and other firms.[13] The original classification for the US is based on a combination of international trade data, geographical concentration as measured by the geographical Herfindahl index and an intuitive sense of which industries respond most to local demand. An equivalent definition in the UK requires a mapping from the four-digit North American Industrial Classification System (NAICS) to the two-digit Standard Industrial Classification (SIC) system used in most of Europe. While the exact mapping between the two systems at different points in time at a very granular level is far from straightforward, the relatively high-level categories used for this analysis in the time period of interest match up almost exactly as can be seen in the first table in Appendix 7.3.[14] The 26 locally non-tradable and 82 tradable four-digit NAICS-12 industries therefore map into 14 and 71 SIC-03 groups, respectively, for this study.

Data from the National Transport Survey (see Department for Transport (2016)) suggest that demand for locally non-tradable purchases is relatively tightly defined. Across England, the average shopping excursion is 7km and makes up a fifth of total trips, though average journeys are two thirds shorter in London and presumably substantially longer in more rural parts of the country.[15] To reflect these patterns, the baseline specification defines locally non-tradable establishment neighbourhoods as a circle with radius 10km. In the robustness tests I explore alternative neighbourhood definitions.

3 Research Design

My empirical strategy tests to what extent monetary policy easing supports consumption and its knock-on effect to local employment.[16] Between autumn 2008 and early spring 2009 the Bank of England cut its policy interest rate from 4.5% to 0.5%. Such an unprecedented monetary easing mitigated the shock to consumption and employment from the ongoing decline in economic activity and house prices. But even if all households benefitted from the support to asset values and their net wealth, only those households on variable-rate mortgages benefitted from the immediate decrease in mortgage payments.[17] This cash-flow effect was dramatic and we would

[13] Mian and Sufi (2014) further subdivide *other* into *construction* and *non-construction* firms.

[14] For example, the NAICS 2012 category of *Automobile dealers* (4411) maps into the SIC 2003 group of *Sales of motor vehicles* (501).

[15] This compares with the average commute within England for work purposes of a little more than 32km. Distances have been increasing marginally over time as household access to cars has increased but there is little evidence to suggest that the number of shopping trips has been falling in any meaningful way, despite the rise of internet transactions.

[16] By using spatial variation, this work can be thought of as the monetary policy analogue to work such as Nakamura and Steinsson (2014) and Dupor and Guerrero (2017).

[17] Byrne, Kelly, and O'Toole (2017) exploit the difference between SVR mortgages and (policy rate) tracker mortgages for the Irish mortgage market. In the UK the policy interest rate also fell further than the SVR but the difference was less sharp.

therefore expect areas with a larger proportion of households on variable-rate mortgages to spend more on local goods and services during 2009 relative to areas with a large number of fixed-rate mortgages. This relative difference in consumption should have translated into a relative difference in employment at these firms to the extent firms adjust their labour in response to demand shocks.[18]

The heart of this study uses the geographic variation in variable-rate mortgagors at the end of 2008 to explain subsequent changes in locally non-tradable employment. Specifically, as shown in the schematic in Figure 3, I estimate the stock of mortgages as of July 2008. Bank rate was initially reduced in October 2008 but I take the stock further back to the summer to exclude those who chose their mortgage type based on the unfolding adverse economic conditions. According to the Bank of England's *Public Attitudes to Inflation* survey, in August 2008 only 10% of people thought interest rates were likely to fall over the next twelve months.[19] I then model the loan-level cash flows from 2008Q3 to 2009Q4 and define the cash-flow shock to be the change in mortgage payments relative to the counterfactual payments had interest rates not fallen.[20] Finally, I compare the proportion of variable-rate mortgages and the neighbourhood cash-flow shocks to the change in establishment-level locally non-tradable employment in the next financial year, between April 2009 and April 2010.[21]

Previous studies such as Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru, and Yao (2017) have aggregated mortgage characteristics and changes in employment to relatively large contiguous administrative boundaries such as zip codes or counties. My preferred specification creates overlapping circular neighbourhoods with centres that lie on the vertices of a 500m x 500m grid that spans the UK. I then assign all establishments to their nearest neighbourhood node and sum up employment across years and nodes. Simplifying to this smaller set of locations brings the number of overlapping neighbourhoods down from around 450,000 to around 50,000 without much loss of spatial accuracy.[22] See Figures 4 and 5 for an illustrative example. Many neighbourhoods only have one locally non-tradable establishment (e.g. a village shop) but out-of-town shopping malls and urban retail districts often have multiple establishments at each location, which means some of the locations are recorded as employing thousands of people. In the main set of results I drop the largest 0.5% of nodes because they are very sensitive to the weighting

[18] Of course, another form of adjustment is to shut down entirely or to start a new establishment in areas where there is judged to be sufficient demand. This is explored in the robustness checks.

[19] In November 2008 that figure had risen to 39%. See question 6 in the Inflation Attitudes Survey at http://www.bankofengland.co.uk/research/Pages/datasets/default.aspx.

[20] This window makes most sense in terms of when spending could affect the employment decisions I study. The results are quantitatively similar for different length windows.

[21] See Appendix 7.7 for robustness checks regarding the timing of the employment response to the cash-flow shock. To the best of the ONS' knowledge, the vast majority of the annual employment data match closely with the financial-year dates reported.

[22] This means that the centre of the closest neighbourhood is at most around 350m away from the true location of an establishment.

Figure 3: Empirical Strategy Timeline

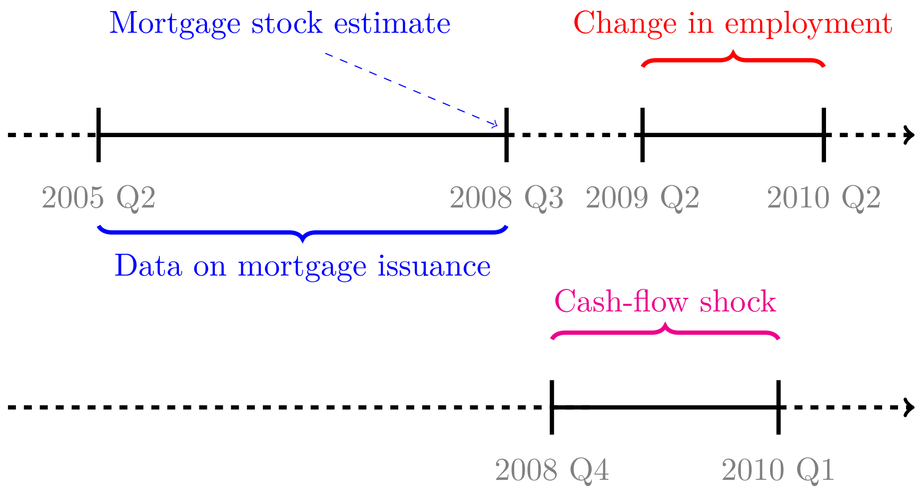

The figure shows the timeline for the central specification. I construct an estimate of the stock in the early autumn 2008 then calculate the income shock associated with lower mortgage payments for every mortgagor over the following five quarters. The dependent variable is the percentage change in employment at the establishment level between April 2009 and April 2010. Since the employment data are from tax records, the staff hiring decision is made a few months before April each year.

but I show the results are almost unchanged in the robustness checks when including the full sample. Individual establishments yet to be born in 2009 or those that had already died by 2010 are assigned zero employees in the employment summation. For example, if one restaurant fails in 2009 but another identical restaurant is created in 2010, employment change at that node is recorded as zero. In the robustness checks I also isolate the intensive margin by only keeping establishments that employed a non-zero number of people in both years.

The fact that my regressions are at the establishment level means that I often have to deal with very large percentage changes (for example, a shop reducing its number of employees from 10 to 1). I therefore want to avoid using log differences, which is a bad approximation to percentage growth for large (negative) changes. In addition, it is helpful to use a measure that is bounded above and below to allow for the possibility of firms shutting down entirely. Following Davis, Haltiwanger, Jarmin, and Miranda (2007), the main outcome variable I use for my analysis is e_i^{10-09}, which represents the percentage change in locally non-tradable employment at establishment i between April 2009 and April 2010.[23]

$$e_i^{10-09} = 2 \times \frac{E_{i,2010}^{NT} - E_{i,2009}^{NT}}{E_{i,2010}^{NT} + E_{i,2009}^{NT}} \tag{1}$$

The initial specification is then

[23] The exact definition of e_i^{10-09} is relatively unimportant. See Appendix 7.8 for robustness checks using the alternative growth-rate definitions.

Figure 4: Assigning Establishments to Nodes

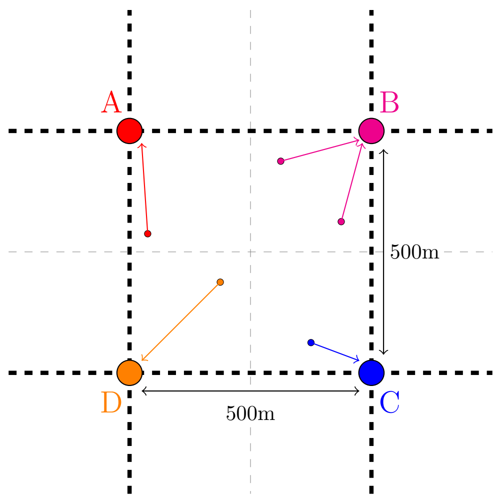

This figure shows how I collapse the 450,000 establishments to around 50,000 nodes without much loss of spatial accuracy. A, B, C and D show nodes on a grid that spans the UK. Establishments are assigned to their closest node and, in this example, two establishments are assigned to Node B. Each node is surrounded by a neighbourhood of customers. Outside densely populated cities, very few nodes contain more than one establishment.

$$e_i^{10-09} = \alpha + \beta \times V_j + \gamma \times X_j + \varepsilon_i \tag{2}$$

The main explanatory variable is V_j, which captures the proportion of mortgaged households opting for variable-rate contracts in neighbourhood j surrounding the establishment i at its centre. The parameter of interest is β, which captures how the proportion of variable-rate mortgagors affects locally non-tradable employment. Unfortunately, it is impossible to know how far customers travel to any given locally non-tradable establishment, let alone generalise for the universe of locally non-tradable firms. I therefore rely on testing different specifications and use transport surveys as a baseline. In the main specification, j defines a circular catchment area with a radius of 10km but alternatives are considered in the robustness checks. These represent an area roughly an eighth of the size of the average US county. A sample of nodes and neighbourhoods is shown in Figure 6. The alternative approach using contiguous administrative regions arguably faces a more severe problem with establishments located near boundaries that are arbitrarily assigned customers that might live far away and are unlikely to visit, while simultaneously excluding nearby customers in the neighbouring region.

Figure 5: Nodes and Neighbourhoods

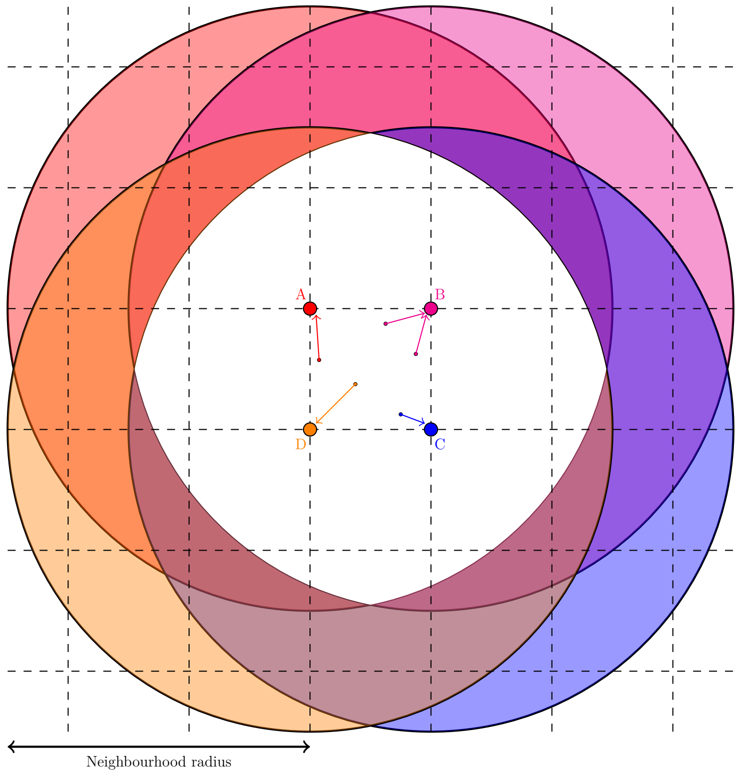

This figure shows the neighbourhoods associated with each example node from Figure 4. It emphasises that although establishments are only assigned to one node, the customer neighbourhoods overlap. In particular, everyone who lives in the white central area is within all four example node neighbourhoods. These neighbourhoods have a 1.25km radius for purposes of illustration.

Figure 6: Oxford Sample Neighbourhoods

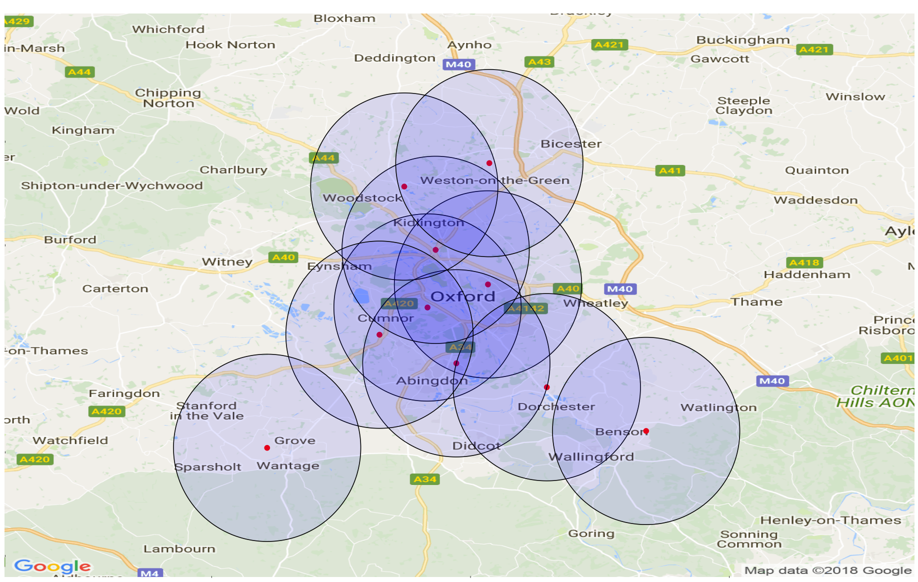

Source: ONS and own calculations. This figure shows a selection of locally non-tradable nodes and neighbourhoods near Oxford assuming the baseline 10km radius.

X_j is a vector of neighbourhood-specific controls including average age, income, loan-to-value (LTV) ratio, loan-to-income (LTI) ratio, house price and net worth change.[24] These controls use the estimated stock of mortgages, so accurately capture the characteristics of mortgagors in each ring. For more information, see Table 3. The combination of geographic accuracy and near-complete coverage of mortgage data for the regression controls is vital for the causal interpretation of β.

Because I model the likely trajectory of every mortgagor's cash flows, I am also able to calculate the average cash-flow shock for each neighbourhood. [25] In Equation 3 the quarterly cash-flow shock is defined as the difference between the modelled quarterly mortgage payment following the fall in Bank rate, $p_{m,t}$, and the modelled quarterly payment assuming Bank rate had not changed, $\tilde{p}_{m,t}$, for each mortgagor, m, at time, t. The average sterling-amount cash-flow shock for neighbourhood j is the sum of these between the start of 2008Q4 and the end of 2009Q4, averaged across all mortgagors in the neighbourhood.[26]

$$\overline{Y}_j^{£} = \frac{1}{M_j} \sum_{m}^{M_j} \sum_{t=2008Q4}^{T} p_{m,t} - \tilde{p}_{m,t} \quad \forall m \in j \tag{3}$$

In Equation 4 the cash-flow shock is scaled by mortgagor annual gross income as of July 2008.[27]

$$\overline{Y}_j^{\%} = \frac{1}{M_j} \sum_{m}^{M_j} \sum_{t=2008Q4}^{T} \frac{p_{m,t} - \tilde{p}_{m,t}}{Y_{m,t_1}} \quad \forall m \in j \tag{4}$$

I can then run a specification that follows the spirit of Mian and Sufi (2014).

$$e_i^{10-09} = \alpha^{\pi} + \delta^{\pi} \times \overline{Y}_j^{\pi} + \theta \times X_j + \varepsilon_i \quad \forall \pi \in \{£, \%\} \tag{5}$$

Here, δ yields the change in locally non-tradable employment growth due to the cash-flow shock. A 1pp increase in the cash-flow shock as a proportion of gross income is associated with a $\delta^{\%}$pp increase in locally non-tradable employment growth.[28] A £1,000 cash-flow shock is

[24] All controls are at the neighbourhood level apart from house price levels and changes, which are at the local-authority level. The change is house prices and GVA is measured between mid 2008 and the end of 2009.

[25] This is approximately equal to the product of the proportion of variable-rate mortgagors and the average shock they experienced. As noted above, however, in practice some fixed-rate mortgagors received a modest income shock as they moved onto the SVR or refinanced between 2008Q4 and the end of 2009.

[26] This includes both principal and interest payments. Those mortgagors on interest-only, variable-rate contracts received the most substantial cash-flow shocks.

[27] I do this by uplifting income-at-origination by average wage growth between mortgage issuance and July 2008. The denominator is adjusted to take account of the fact five quarters of payments are summed but I only have data on annual income.

[28] I refer to $\delta^{\%}$ as the main semi-elasticity of interest. Specifically, it is defined as $\frac{df(x)}{dx}x$, so a percentage change in the independent variable (cash-flow shock, as a proportion of income) is associated with an absolute change in the dependent variable (annual growth rate of locally non-tradable employment).

associated with a $\delta^{£}$pp increase in locally non-tradable employment growth. All regressions are weighted by node employment.

3.1 Identification

A key identification assumption that spans this whole research design is that households with variable-rate mortgages increased their relative consumption because they received a cash-flow shock rather than because they differed from fixed-rate mortgagors. The ideal experiment involves randomly assigning households different amounts of money and observing their behaviour. Such helicopter drops are not found in the real world and so I exploit quasi-random variation in the UK mortgage market for my identification strategy. I need to demonstrate that the effect I find is not biased by selection of mortgagors into different types of mortgage ex ante or leakage between different types of mortgage ex post.

3.1.1 Selection Bias

The first concern is that certain types of households selected into fixed and variable-rate mortgages between 2005 and 2008. If more financially secure households opted for variable-rate mortgages this might lead to downward bias in my estimate of β and δ because some richer households might have a lower marginal propensity to consume.[29] That said, the direction of bias is unclear given the recent work on high marginal propensities of consumption for the *wealthy hand-to-mouth* (e.g., Violante, Kaplan, and Weidner (2014)). Either way, I argue there was very limited selection along the interest-rate fixation dimension and the observable characteristics between those on fixed and variable-rate mortgages were remarkably similar.

Mortgages are poorly understood by households and each mortgage contract has a vast number of non-price terms and options.[30] The choice between a fixed and a variable rate is much less important in the UK than in other parts of the world because the length of most fixed-rate contracts is a few years at most. The 75% LTV two-year variable and 75% LTV two-year fixed-rate mortgages have been consistently the most popular mortgage products for the last twenty years. Consequently, during the Great Moderation when interest rate volatility was modest, the lifetime cost of choosing a variable-rate contract and timing it poorly was on the order of only a few thousand pounds (or, a couple of percentage points of house value). This stands in stark contrast to the US market where the mortgage choice was often between a 30-year fixed rate with the option to refinance and an adjustable-rate mortgage with a tempting teaser rate. Due to the perception of the relative unimportance of the evolution of the interest rate, many UK household

[29] To the extent wealthier households are often better able to tolerate financial risk, this would be consistent with the theoretical findings of Campbell and Cocco (2003).

[30] Agarwal, Chomsisengphet, Liu, and Souleles (2015) shows people struggle to calculate the present value of fees compared to additional percentage points on the underlying interest rate.

decisions therefore put much more weight on other parts of the contract, such as whether the principal could be pre-paid; whether the loan could be ported to a new property; or how much the lender was willing to lend relative to collateral or income.

Identification would be cleanest if the choice of a fixed or variable-rate mortgage was truly random for an individual. If that were true, and each person had around a 50% chance of choosing a fixed-rate contract, we would expect to see just under half of all people switching to a different contract the next time it came up for renewal.[31] I use the population of remortgagors in my sample to explore this. The top panel of Table 1 shows the expected distribution of mortgage choices under the null hypothesis that the probability of choosing a fixed-rate mortgage was equal to the sample average and independent across time. The bottom panel shows the observed behaviour of remortgagors. It turns out that over 37% of households that remortgaged between 2005 and 2008 moved onto the opposite interest-rate contract, only 10pp lower than expected under the null hypothesis.

Not only do a surprisingly large share of households flip mortgage type but the switch is relatively symmetric. This evidence suggests that the type of interest-rate contract households chose did not give a good signal about their characteristics or preferences. Although the choice is unlikely to be truly random, I argue it is such a marginal decision for many people that it provides the quasi-random variation that I can exploit. The instrument I develop later on uses the fact that, even if there are time-varying factors that explain mortgage interest rate decisions, the timing of mortgaging events is also quasi-random. This is especially true in the UK given that over half of all mortgages are remortgages and the vast majority of people refinance at very regular intervals (when their deals expire).

The quasi-random assignment is also consistent with the observable characteristics of the mortgagors reported in Table 2, which look broadly similar across mortgagor types. The median income and house price is, however, slightly higher for variable-rate mortgagors. Given the median loan is similar across both mortgagor types, it follows that those on variable rates on average have a lower LTV and LTI ratio, which are common proxies for vulnerability.[32] The latter, especially, might be cause for concern if it means that those on fixed rates are likely to be more sensitive to counterfactual cash-flow shocks.[33]

Although the median values in Table 2 suggest that there might be some inherent difference between the types of people opting for different mortgages, there a four grounds for reassurance.

[31] Since $2x(1-x) \approx 0.5$ for $x \approx 0.5$.

[32] Bacon and Moffatt (2012) find that those on fixed-rate contracts were likely to borrow more between 1992 and 2001, which is not the case in the PSD data.

[33] The PSD data also show that those issued with variable-rate mortgages were less likely to pay back their principal (i.e. have a so-called interest-only mortgage). This is consistent with the findings from Piskorski and Tchistyi (2010) who show that interest-only mortgages might benefit those looking to buy expensive houses on leverage with variable incomes.

Table 1: Evidence for Quasi-Random Mortgage Selection

Expected | no intertemporal relationship, percent

	Second mortgage	
First mortgage	Variable	Fixed
Variable	13.7	23.3
Fixed	23.3	39.7

Note: Pr*(Fixed) = 63%*

Observed transition matrix, percent

	Second mortgage	
First mortgage	Variable	Fixed
Variable	18.5	16.8
Fixed	20.1	44.6

There were 763,276 people who remortgaged once between 2005 and 2008 and, in total, 63% of the mortgages were fixed-rate contracts. If the probability of choosing a fixed-rate contract was independent across time we would expect to see the distribution in the top panel. For example, the probability of choosing a fixed-rate contract twice is $0.63 \times 0.63 = 39.7\%$. The observed distribution is shown in the bottom panel.

First, post-2005 mortgagors as a whole are a relatively homogeneous group of people. Unlike outright owners or renters, they are overwhelmingly likely to be aged between 25 and 60, have stable incomes and have similar consumption baskets. Differences between fixed and variable rate-mortgagors are small compared to those between mortgagors and other tenure groups.

Second, discretionary spending habits tend to be non-linear and so a comparison of those in the bottom quartile (or top quartile of LTI) is likely to be the more informative than median differences: these points of the distributions are more similar (second and third columns of Table 2). Focusing on the difference between the median or the mean is likely to overstate the behavioural heterogeneity of the two groups.[34]

Third, the observable characteristics are based on mortgagors at origination. Since quite a few fixed-rate mortgagors rolled off their contract onto a variable rate and received the cash-flow shock, any segregation of type at origination is diluted by the time interest rate fall at the end of 2008.[35] Finally, to the extent that the poorest mortgagors are likely to prefer fixed-rate

[34] Bunn and Rostom (2015) suggest there might be non-linearities in behaviour due to income gearing. For example, in the Crisis, the significant falls in consumption occurred for households with LTIs greater than 4 so on this metric the mortgagor behaviour is likely to be similar across the interest rate types. The same work shows that consumption responses of renters and owner occupiers were small compared to those of mortgagors with an LTI greater than 2. For this reason I only focus on the consumption response of mortgagors in this study. It also helps to justify why the pre-2005 mortgages that are missing from the PSD data set are likely not to be important since they are likely to have lower LTIs. See Flodén, Kilström, Sigurdsson, and Vestman (2017) for more discussion on the role of indebtedness and the mortgage cash-flow effect.

[35] For the most part, the switch from fixed to variable rates would have been determined by timing in a similar way to in Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru, and Yao (2017) and Flodén, Kilström, Sigurdsson,

mortgagors across time, this makes Table 1 even more convincing. Despite this possible selection bias for the most extreme households in the distribution, the empirical transition probabilities are still very close to those expected under the null hypothesis. These transition probabilities are explored more in Appendix 7.4.

It is also worth noting that identification is only compromised to the extent there are unobservable differences between those on fixed and variable interest rates. Given the controls I employ are both extensive and accurate, I argue that the main specifications yields the causal effect of cash-flow shocks on locally non-tradable employment.[36] For example, although I have not explicitly controlled for risk aversion, this is likely to be correlated with observable characteristics such as income, age and house price.

3.1.2 Mortgage Type Leakage

The second concern with my identification strategy is that some households might switch between fixed and variable-rate mortgages after the fall in interest rates. This could lead to upward bias in my estimate of β if those wanting to increase consumption were more likely to take advantage of the financially beneficial mortgage switch.

The active leakage described above is actually unlikely to have much of a bearing on my results. Following the collapse of Lehman Brothers, the UK mortgage market became severely impaired. As collateral values fell, mortgage lenders withdrew from the once-active remortgage market. Figure 7 shows how the number of remortgages per month fell by around 75% by the start of 2009. Even if all remortgages over the next twelve months were fixed-rate mortgagors keen to refinance onto an interest rate closer to Bank rate, that only accounts for around 7% of the stock of mortgages used in the main regressions.[37] This subdued refinancing activity stands in stark contrast to the US where remortgages spiked up following the fall in interest rates as those on long term fixed-rate mortgage contracts took advantage of the lower rates.

and Vestman (2017). During the Great Recession some mortgagors moved off fixed-contracts onto variable rates because their collateral had fallen in value limiting their options. This is not a cause for concern for the mortgage stock in 2008 because house prices had been increasing almost universally across the country.

[36] That is, accurate for post-2005 mortgagors.

[37] There were also around 400,000 home movers during this period.

Table 2: Mortgagor Statistics in July 2008

At Origination	Median	25^{th}	75^{th}	Mean	N
Variable-rate					
Interest only				0.35	1,917,824
Verified income				0.53	1,917,824
Joint mortgagees				0.50	1,917,824
Single income	35,484	24,813	56,000	56,193	960,174
Joint income	49,500	35,857	72,000	67,227	957,650
Age	40	33	47	40.8	1,917,824
Loan value	110,000	70,000	169,775	140,132	1,917,824
House price	194,000	137,000	285,000	248,812	1,917,824
Mortgage term	20	15	25	19.9	1,917,824
LTV	61.9	40.4	81.0	59.9	1,917,824
LTI	2.6	1.9	3.4	2.67	1,917,824
Fixed-rate					
Interest only				0.26	3,845,703
Verified income				0.56	3,845,703
Joint mortgagees				0.53	3,845,703
Single income	30,000	22,172	44,040	40,826	1,808,014
Joint income	43,348	33,350	58,666	53,072	2,037,689
Age	36	30	44	37.5	3,845,703
Loan value	109,316	75,000	155,000	127,771	3,845,703
House price	165,000	120,000	230,000	198,132	3,845,703
Mortgage term	24	19	25	22.6	3,845,703
LTV	73.0	51.8	88.1	68.1	3,845,703
LTI	3.0	2.3	3.6	2.96	3,845,703

This summary table shows the observable characteristics of mortgages at the end of 2008Q2. Interest-rate types are those designated at origination, which means in practice many of the fixed-rate mortgagors had moved on to the SVR by the summer of 2008.

Figure 7: Mortgage Approvals

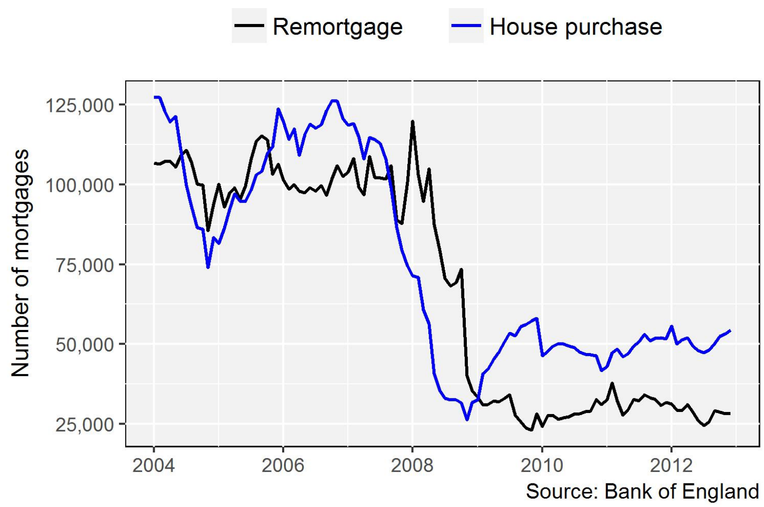

The series correspond to the seasonally adjusted data for housing remortgage (B4B3) and housing purchase (VTVX) from the Bank of England Statistical Interactive Database.

The low refinancing activity in the UK is partly a result of weakened credit supply conditions but also those of credit demand. High early repayment fees in the UK disincentivised those on fixed-rate mortgages to break their contract early. Moreover, lower collateral values and suppressed lender tolerance for high-LTV mortgages meant that refinancing interest rates stayed stubbornly high for many households. Between the second half of 2008 and the end of 2009 the mean recorded remortgage rate for people previously on fixed-rate deals was 4.9%, only around 70bp below the mean fixed-rate interest rate on the stock. Indeed, the SVR proved an attractive alternative to remortgaging, especially in cases where collateral values had fallen.

Relatively short fixation periods meant that most households on fixed-rate mortgages preferred to wait for their contract to roll off and move on to the SVR to benefit from lower interest rates. This passive form of leakage is less concerning for my results. My β coefficient takes account of the decision to opt for a fixed or variable-rate mortgage before the policy easing; before there is any incentive to switch. The gradation in the cash-flow shock that results from fixed-rate contracts expiring at different times is captured in my estimates for δ.

4 Results

4.1 Summary Statistics

The combination of the mortgage interest-rate type and leverage diffracted the direct effect of monetary policy into a heterogeneous cash-flow shock across mortgagors. Figure 8 shows the distribution of that cash-flow shock across mortgagors on both fixed and variable-rate contracts.

The majority of people on fixed-rate contracts saw no change in mortgage payments when interest rates fell so the modal bar has been removed for ease of comparison between the charts; just over half of mortgagors on fixed-rate contracts received a cash-flow shock of less than 0.5% of annual gross income between September 2008 and the end of 2009. Nevertheless, because the average length of the fixed period was a little over two years, many fixed-rate households were able to roll onto the SVR at some point and receive a partial cash-flow shock, so the mean shock was 0.9%.

Figure 8: Individual Cash-Flow Shock Distributions

(a) Fixed-Rate Mortgagors **(b)** Variable-Rate Mortgagors

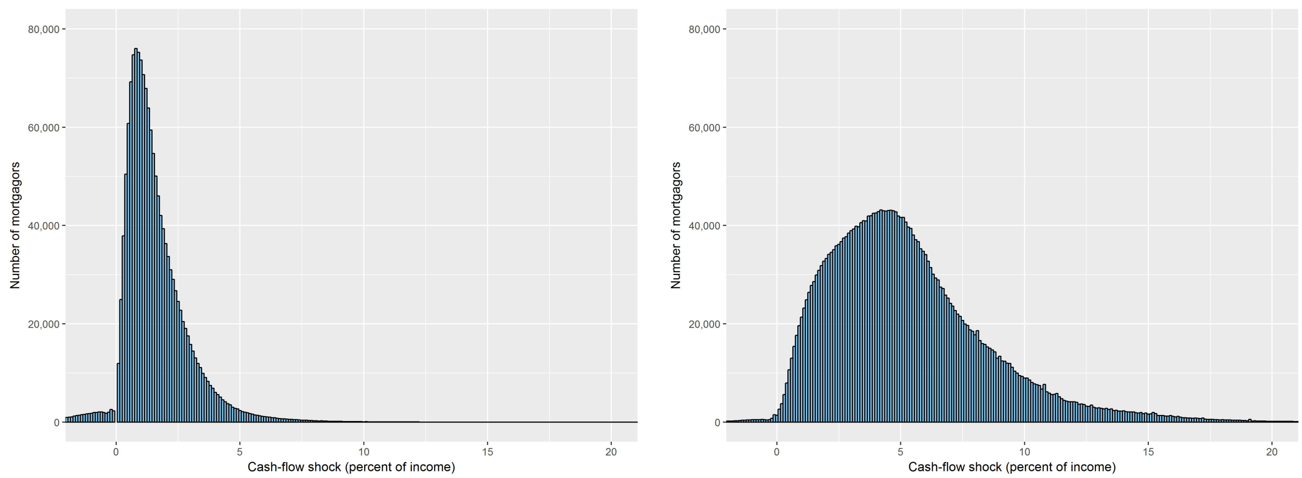

Source: PSD. A total of 4.5 million mortgages are captured in these distributions. 1.2 million mortgages very close to 0 have been dropped from the fixed-rate distribution to make the charts easier to compare. The small number of negative cash-flow shocks occur when people revert or remortgage onto a less favourable rates.

Those on variable-rate contracts received a mean cash-flow shock of just over 5% of gross income. This is an economically significant number; only slightly below the household saving rate during the Great Moderation. For many, it was equivalent to around 8% of post-tax income and perhaps closer to 30% of annual discretionary income (after subtracting food, travel, etc).[38] This cash-flow shock was substantially more for those who had borrowed more against their income or those on non-amortising mortgages; around 8% of variable-rate mortgagors received a cash-flow shock of more than 10% of gross income. The spatial distribution of both types of mortgagor led to the neighbourhood income shock distributions shown in Figure 9.

[38] See ONS data series NRJS. This is also consistent with survey evidence. For example, households saved around 8% of income in 2013 (Bunn, Rostom, Domit, Worrow, and Piscitelli (2013)). Median household disposable income per head in 2010 was around £16,000.

Figure 9: Neighbourhood Cash-Flow Shock Distribution

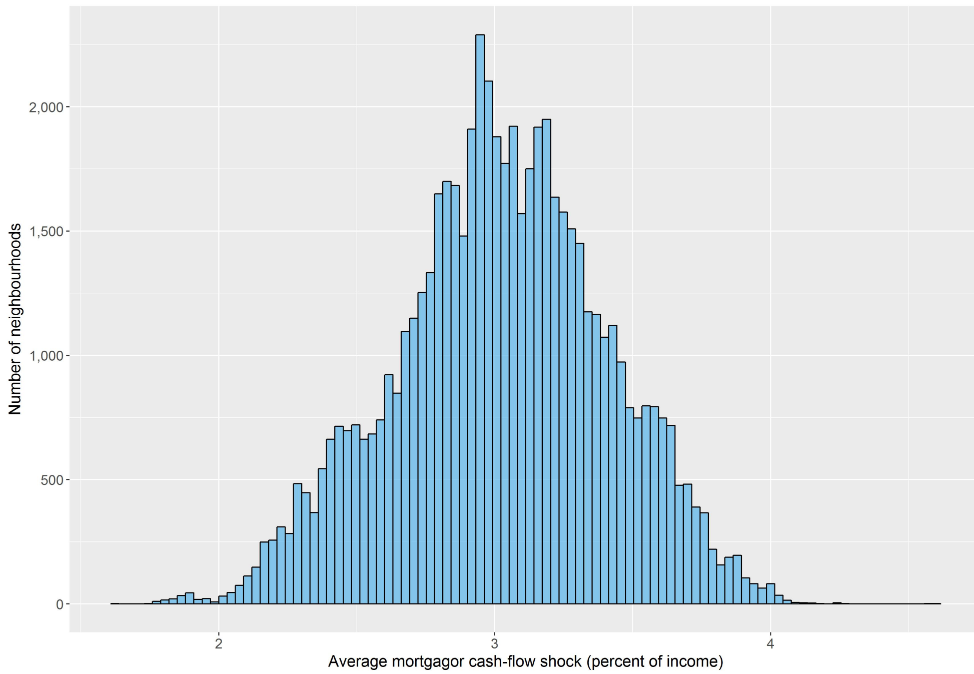

Source: PSD. The average cash-flow shock for mortgagors in a neighbourhood is a function of, among other things, the proportion of people entering fixed and variable-rate contracts, average household income gearing and the when the contracts were taken out.

Table 3 presents the summary statistics associated with the main set of regressions. It shows that the main sample consists of almost 50,000 overlapping neighbourhoods, each containing an average of just under 20,000 mortgages. Just over half of mortgaged households received a substantial favourable cash-flow shock though only a third of households had entered a variable-rate contract at origination. The average neighbourhood mortgagor cash-flow shock was about £2,000, or 3% of gross income. Most of the neighbourhoods had a very small establishment (or group of establishments) at their centre. But the distribution of neighbourhood employment is heavily positively skewed, as shown by the fact the mean employment is more than double the upper quartile. The (unweighted) average neighbourhood experienced a 17% reduction in locally non-tradable employment.

4.2 Aggregate Cash-Flow Effect

In this paper I argue that an important determinant of employment growth is the cash-flow channel of monetary policy. Those who experienced substantial falls in mortgage payments were able to continue spending on local goods and services. Relative spending increases were associated with relative differences in annual employment growth rates. My first result is to show that locally non-tradable establishments surrounded by a large proportion of variable-rate mortgagors experienced a relative increase in employment. The first column of Table 4 shows

Table 3: Neighbourhood Summary Statistics

Statistic	N	Mean	St. Dev.	25^{th}	75^{th}
Share of flexible-rate mortgages, %	47,261	53.6	5.34	50.3	55.2
Share of flexible at origination, %	47,261	35.3	6.94	30.8	37.0
Average cash-flow shock, £000	47,261	1.97	0.70	1.47	2.23
Average cash-flow shock, %	47,261	3.02	0.39	2.77	3.28
House price 2008, £000	47,261	216	65.6	167	247
Change in net worth, 2008-2009, %	47,261	18.1	17.5	10.1	29.7
Loan value, £000	47,261	128	43.6	97.8	147
LTV	47,261	63.1	4.36	60.0	66.4
LTI	47,261	2.85	0.17	2.75	2.98
Age	47,261	39.3	1.77	38.0	40.4
Single income, £000	47,261	44.6	13.5	35.2	49.7
Joint income, £000	47,261	56.4	13.4	47.2	61.1
Number of mortgages, 000	47,261	18.8	29.0	3.52	22.7
Number of households, 000	47,261	75.0	124	14.6	87.4
Population, 000	47,261	182	299	34.9	209
Construction employment share, %	47,261	5.98	2.54	4.48	6.86
Weighted GVA level, £000	47,261	43.5	7.34	38.7	47.1
Weighted change in GVA, 2008-2009, %	47,261	-2.59	-2.04	-3.35	-1.57
Employed proportion, %	47,261	42.0	12.4	33.6	48.3
LNT employment (at node), 2009	47,261	60.7	222	3	26
Change in LNT employment (at node), 2009-10, %	47,261	-16.6	65.8	-22.2	6.45

This table presents summary statistics for the neighbourhood data used in the analysis. The final two columns refer to percentiles. House prices are from the ONS local-authority series. The number of households and population estimates are constructed using postcode-level census data. Locally non-tradable (LNT) employment data are from the BSD and the neighbourhood GVA shocks are constructed using the neighbourhood share of employment in each of the 17 main industrial categories. All other data are constructed from the PSD. These overlapping neighbourhoods represent more than 99% of total UK population and LNT employment.

my estimate for β. It suggests that a 10pp, or 2 standard deviation, increase in the proportion of variable-rate mortgagors increased employment growth in the locally non-tradable sector by about 1.9pp over the next year. As well as quantifying the cash-flow effect of monetary policy, this also provides an insight into how monetary policy pass-through is affected by the structure of the mortgage market.

The second two columns of Table 4 present estimates for δ in Equation 5 using definitions in Equations 3 and 4. Column 2 suggests that a cash-flow shock equivalent to 1% of gross income was associated with an increase in locally non-tradable employment growth of a little over 4.5pp and is statistically significant at the 1-percent level when standard errors are clustered by local authority (the choice of clustering standard errors at the local authority is conservative, see Appendix 7.9 for further details). Column 3 suggests that a cash-flow shock equivalent to £1,000 is associated with an increase in locally non-tradable employment growth of around 7pp. Full details regression details can be found in Table 17 of Appendix 7.5.

Using the coefficient in the second column and the fact that a 400bp fall in policy interest rates yielded an average mortgagor cash-flow shock of 3.0%, my results suggest that a 100bp accommodative monetary policy shock led to locally non-tradable employment growth of just over 3.5pp the following year through the cash-flow channel.[39] That is a large number but plausible given there are good reasons to think that the monetary policy shock identified in this study was perceived by households to be permanent. This large coefficient implies the counterfactual locally non-tradable employment growth would have been markedly weaker, and therefore unemployment markedly higher, if monetary policy had been even slightly less aggressive during the Crisis. Between 2009 and 2010, locally non-tradable employment shrank by around 3%. Taken at face value, my results therefore suggest locally non-tradable employment could have shrunk by more than 15 percent absent policy intervention. Such calculations should, however, be treated with caution because of the partial equilibrium nature of the cash-flow shock channel identified.

4.3 Regional Heterogeneity

The diffraction of the cash-flow shock through the regional mortgage and labour market structures led to a heterogeneous spatial impact of monetary policy easing. In the UK the most natural administrative region to illustrate my results is the local authority. Taking the full 400bp of policy easing, Figure 10 shows the estimated impact of the accommodative monetary policy shock on all 389 local authorities' total employment growth by quartile.[40] Areas shaded dark (light) green correspond to parts of the country where overall employment was most (least) affected by

[39] Using the sterling-amount coefficient in the third column, the analogous calculation is that a 400bp fall in policy rates led to an average cash-flow shock of around £2,0000 and an employment effect of 3.4pp. See Table 3.

[40] The Isles of Scilly are excluded and the City of London is subsumed into the City of Westminster for purposes of clarity.

Table 4: Results Summary

	Annual employment growth, pp		
	(1)	(2)	(3)
Variable rate (% mortgagors)	0.19***		
	(0.06)		
Cash-flow shock (% income)		4.65***	
		(1.44)	
Cash-flow shock (£000)			6.81***
			(2.55)
Controls	**yes**	**yes**	**yes**
Observations	47,261	47,261	47,261
Note:			*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level. Neighbourhood controls include average house price, change in household net worth, loan value, age, income, change in neighbourhood GVA and LTI. Only the latter is statistically significant.

monetary policy easing through the cash-flow channel.[41]

Before exploring Figure 10 in detail, it is worth putting this regional heterogeneity in context. First, the cash-flow channel is an important part of the transmission mechanism but it is far from the only one. Among other things, accommodative monetary policy altered expectations, incentivised people to reallocate resources (across time and purpose) and boosted asset prices. These effects are not captured in my work. Second, one person's cash-flow boost is someone else's cash-flow compression in a closed system. The flip side of lower mortgage repayments was lower income for the banking system and agents providing its funding. Of particular importance here, households received less income on their deposits, which partially offsets the positive cash-flow shock on the liability side of household balance sheets.[42] Finally, there are good reasons to think that the cash-flow effects I find would likely have been approximately symmetric if interest rates had instead increased.

The spatial distribution of the total cash-flow effect is the confluence of a number of factors. In Northern Ireland, monetary policy had a particularly strong effect because there was a large proportion of floating rate mortgages, as shown above in Figure 1. On the other hand, in Cornwall (SW England) the driving factor was the large share of people employed in the locally non-tradable sector. The monetary policy effect is below the median in the vast majority of Scotland due to a combination of a large proportion of fixed-rate contracts and limited borrowing against incomes. So if interest rates had instead increased, areas such as Scotland would have been more insulated from the corresponding cash-flow effects.

[41]See Appendix 7.6 for robustness test on the central specifications by region.

[42]See Cava, Hughson, and Kaplan (2016) for why the overall cash-flow shock was positive for the economy.

For the UK as a whole, mortgage type, income gearing and employment structure all contributed relatively evenly to the regional heterogeneity of the cash-flow effect. This heterogeneity would have been most acute between 2009 and 2010 while mortgagors were transitioning off their (relatively short) fixed-rate contracts. By 2011, most mortgages were linked to policy rates and spending choices had adjusted to the new low-interest-rate environment. In terms of examining how this heterogeneity might evolve during future episodes of monetary policy action, mortgage-type distributions are likely to change relatively quickly, though certain characteristics such as age and income continue to drive mortgaging choices on the margin. Today, much of Scotland still has more fixed-rate mortgagors than the UK average, but some trends are common across all regions. For example, the average duration has increased for all new mortgages. Income gearing is modestly persistent and will remain important for regional variation as long as there continues to be variation in house prices relative to local incomes. Finally, the labour market structure is probably the most slow-moving contributor to regional variation and is likely to be especially important when looking at very local cash-flow effects. The balance between tradable and locally non-tradable employment can be thought of as a microcosm for cross-country differences in employment structures and current-account openness.

Figure 10: Estimated Cash-flow Effect on Total Employment Growth, 2009-10

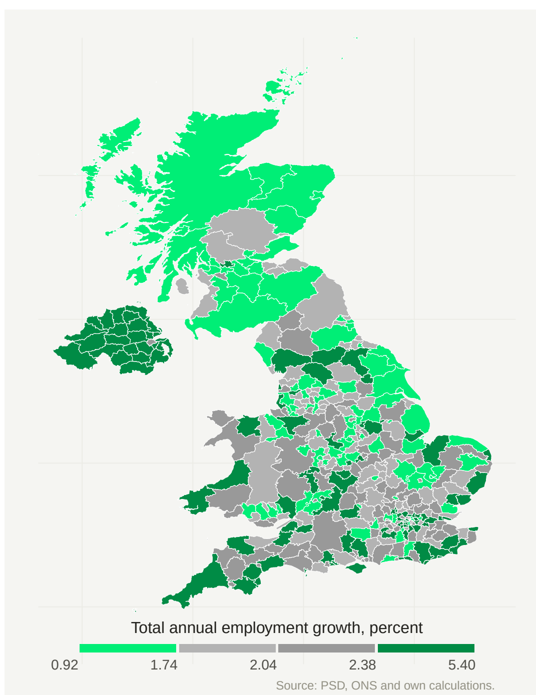

This figure is constructed by taking the establishment-level cash-flow semi-elasticity of employment and combining it with cash-flow shocks and locally non-tradable employment shares for each local authority. Colour breaks denote quartiles.

Table 5 provides more details on the distribution of the estimated overall effect of monetary policy across regions. Column 1 combines $\hat{\delta}^{\%}$ with the share of locally non-tradable employment and cash-flow shock across each local authority to find the regional distribution of monetary policy impact. The difference between the 95th and 5th percentile is 1.55pp and this is taken as my central measure of regional variation. Column 2 treats the distributions of the estimated effect on locally non-tradable employment and the share of that employment separately. I then take the product of these values at each percentile to construct a counterfactual regional variation, which gives a plausible upper bound of how the overall effect of monetary policy might vary. Columns 3 and 4 repeat this exercise for the estimates relying on the absolute cash-flow shock, $\hat{\delta}^{£}$, and yields quantitatively similar dispersion measures.

Between 2008 and 2010 aggregate unemployment only increased by around 3pp so the estimated variation in the effect of monetary policy of around 1.5pp across regions is sizable. But even taking the inter-quartile range as the measure of dispersion shows that regional difference still matter. That 64bp difference translates to almost 500 jobs for the average-sized local authority. Finally, it is worth noting that although local authorities are a useful regional subdivision for illustrative purposes, such aggregation does mask some of the the regional heterogeneity that exists when exploiting the establishment-level identification used in this work. The sub-regional variation in the effect of monetary policy is likely to be stronger still.

5 Robustness

5.1 Instrumenting for Mortgage Choice

Households often paid more attention to the initial interest rate when choosing a mortgage than the set of future contractual obligations, which were often difficult to understand and of limited practical importance in a world of monotonically increasing asset values. The top panel of Figure 11 shows how much variation there was in the origination of fixed and variable-rate mortgages in the five years before the Crisis. At the start of 2005, for example, two thirds of new mortgages were issued with a fixed interest rate. Nine months later that proportion had fallen to around a third. The bottom panel suggests that people found it difficult to compare the present value of the different mortgages on offer: when the fixed-rate premium (FRP) (the spread between the available two-year 75% LTV fixed-rate mortgage and the two-year 75% LTV variable-rate mortgage) increased, fewer people opted to fix. This is consistent with evidence from Miles (2005) and Badarinza, Campbell, and Ramadorai (2017) and the notion that people pay more attention to the headline rate than whether the overall cost of the mortgage is consistent with their expectations regarding the path of policy rates.

Although the FRP captures the market expectation of the future path of interest rates, there are good reasons to think this information does not have much impact upon household choices.

Table 5: Estimated Heterogeneous Effect of Cash-flow Channel

Local Authority Percentile	Annual total employment growth, pp %-shock Observed	Counterfactual	£-shock Observed	Counterfactual
0	0.92	0.72	0.40	0.40
5	1.32	1.07	0.63	0.73
10	1.49	1.28	0.72	0.88
25	1.74	1.68	0.89	1.26
⋮	⋮	⋮	⋮	⋮
75	2.38	2.53	1.57	2.64
90	2.74	2.99	1.97	3.89
95	2.87	3.27	2.35	4.68
100	5.40	5.54	13.05	20.61
75-25 spread	**0.64**	**0.85**	**0.68**	**1.38**
90-10 spread	**1.26**	**1.71**	**1.24**	**3.01**
95-5 spread	**1.55**	**2.20**	**1.71**	**3.96**

This table shows local authority percentiles associated with the a 400bp cut in monetary policy constructed using the establishment-level coefficients. Unlike previous regression tables, *total* refers to all private-sector employment.

Many people struggle to understand the mechanics of mortgage finance and the Bank of England's *Public Attitudes to Inflation* survey over time shows that people do not consider the market curve when deciding how interest rates might move over the next few quarters.[43] Even if households understand how to interpret market prices, it is plausible that many (implicitly) discounted future mortgage payments because they thought their income would increase or higher collateral values would ease the burden of future refinancing.[44] In either case, I can exploit the explanatory power of the FRP in determining mortgage choices by using it as an instrument as long as mortgagors based at least part of their decision on whether to take out a fixed or variable-rate mortgage on the relative price of the initial interest rate.

[43] Evidence presented in Agarwal and Mazumder (2013) suggests low mathematical ability is associated with poor financial decision making. My identification relies more on the level of overall mortgage market comprehension.

[44] See Cloyne, Huber, Ilzetzki, and Kleven (2017) for evidence in the UK.

Figure 11: Mortgage Issuance

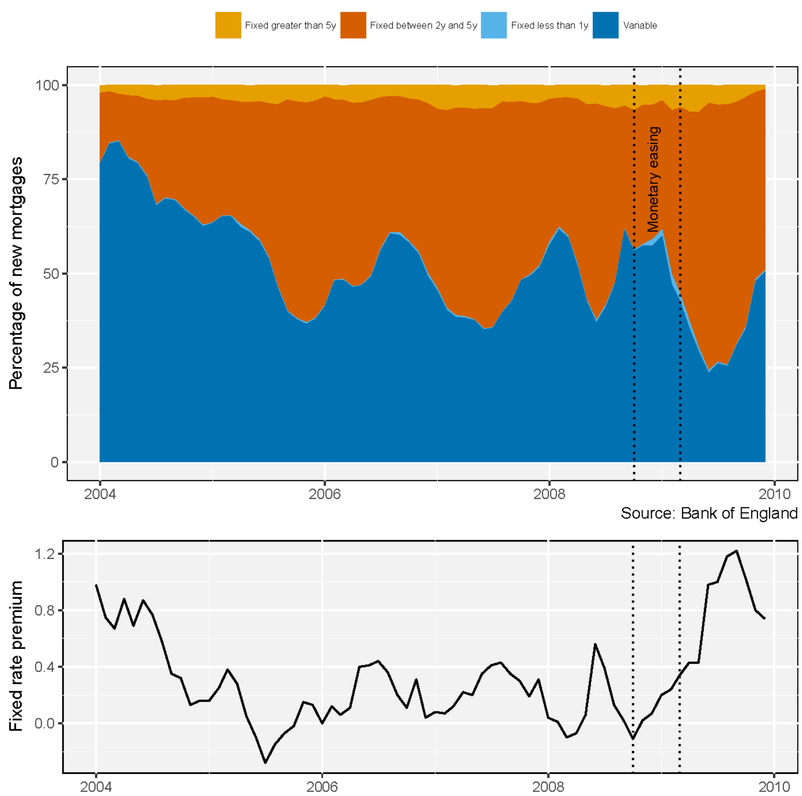

The top panel of this figure shows the variation in mortgage-type issuance across time. The bottom panel shows the time variation in the fixed-rate premium. These data are taken from Bank of England aggregated data and therefore have a slightly different make-up from the rest of the analysis (e.g., they include second-charge mortgages). The series correspond to Fixed for more than 5 years (CFMB9I2 and CFMB9I3), Fixed between 1 and 5 years (CFMB8R8), Fixed less than a year (CFMB8R7) and Variable rate (CFMB8R5). In practice, there is a negligible number of mortgages issued with contractual maturities of between 1 and 2 years.

The UK mortgage market is competitive and there is surprisingly little variation across lenders in the two benchmark rates they offer over time. The large UK lenders have almost no geographical bias in their lending activities so I rely on pure time variation in the FRP, which is driven by changes in the steepness of the short-end yield curve. To construct my instrumental variables estimate I follow a three-step least squares approach.[45] The first stage requires constructing a linear probability model to predict the probability that an individual will take out a variable-rate mortgage, P_m, based on the average FRP prevailing in the three months before the mortgage was issued. X_m represents a host of other individual controls, which align with the main neighbourhood-level regressions, including LTI, LTV, house price and age.[46]

[45] This approach is similar to the one taken in Adams, Almeida, and Ferreira (2009).

[46] The dependent variable is a dummy variable representing whether the individual took out a variable-rate mortgage at origination and not the status of the mortgage in July 2008. Since the model fit is good I do not need to worry about the probability being below 0 or above 1, though a probit model gives very similar estimates at the neighbourhood level.

$$P_m = \alpha + \pi \times FRP_t + \lambda \times X_m + \varepsilon_m \tag{6}$$

The results from this first stage are shown in the first column of Table 6. The central message is that a 10bp increase in the FRP was associated with the probability of taking out a fixed-rate mortgage increasing by around 1.4pp. That suggests between the middle of 2005 and the middle of 2006 the steepening yield curve reduced the probability of someone taking out a fixed-rate mortgage by around 10pp. The five-digit F statistic shows the instrument is relevant to a very high level of significance.

The first stage is conducted at the individual level but the ultimate outcome of interest is the change in employment at locally non-tradable establishments, which rely on all customers in the local neighbourhood for their revenues. I therefore need to aggregate the fitted values obtained in the individual regression to the neighbourhoods in my main specification. This is done by taking a simple average across mortgaged households, m, that belong to neighbourhood j. It yields the expected number of fixed-rate mortgages at origination for each neighbourhood.

$$\hat{\mathcal{P}}_j = \frac{1}{M_j} \sum_{m=1}^{M_j} \hat{P}_m \quad \forall m \in j \tag{7}$$

I am now in a position to use the aggregated fitted values in Equation 7 as an instrument for the proportion of variable-rate mortgages in my neighbourhoods. Figure 12 gives a sense of the instrument's high relevance by plotting neighbourhoods bucketed by actual and predicted number of mortgagors on fixed-rate contracts.

Figure 12: Instrument Relevance

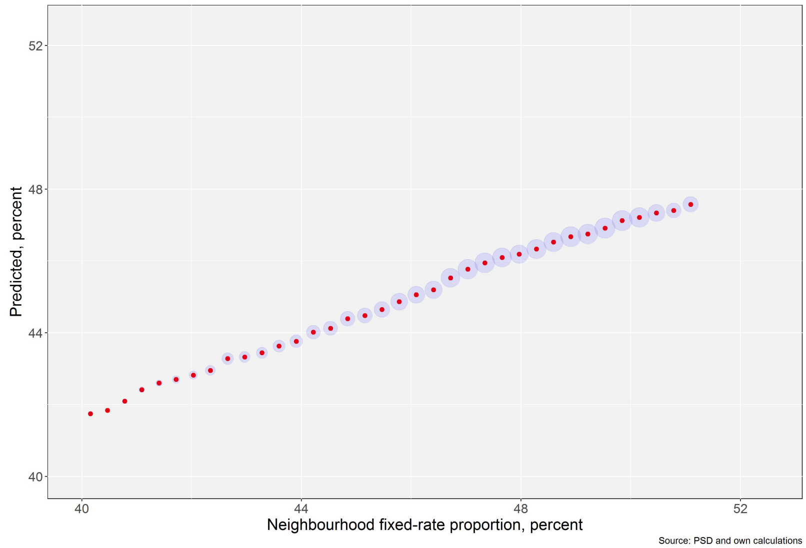

This figure plots the neighbourhood proportion of fixed-rate mortgagors in the summer of 2008 against the predicted proportion of mortgagors taking out fixed-rate contracts after aggregating individual fitted probabilities (in Equation 7). Blue discs represent the number of mortgages in each bucket. Top and bottom 1% of observations have been dropped.

This approach has the advantage that it uses the loan-level information in the first stage but the linear probability specification need not be correctly specified for the general approach to be valid. More importantly, because the fitted values are used as an instrument, the individual and neighbourhood models can use different information. For example, the change in net wealth after 2008 might have directly influenced spending but it will not have directly influenced the previous decision to take out a particular mortgage.

The second column of Table 6 shows the instrumental variables estimate, which is very close to the main coefficient presented in Tables 4 and 17. This provides some reassurance that the main specification has been well identified. One concern with the specification in Equation 2 is that omitted variables could determine both mortgage choice and consumption responses to cash-flow shocks. For example, if those on high incomes were more likely to choose variable-rate mortgages and they had lower marginal propensities for consumption, that would lead to bias downwards in the estimate of β. The results using the FRP as an instrument suggest that either this selection bias is very limited or the observable characteristics of mortgagors do a good job of controlling for it, and therefore the estimate of β can be interpreted as causal.

5.2 Types of Firm

The results in Table 17 suggest that the cash-flow effect of monetary policy has a significant impact on the employment of firms that rely on customers living nearby. To strengthen this result, I investigate whether the same cash-flow shock has a null effect on firms that do not sell

Table 6: Instrumental Variables

	Mortgagor level P(fixed-rate)	**Neighbourhood level** Annual employment growth, pp
Fixed-rate premium (bp)	-0.139^{***}	
	(0.001)	
Variable rate (% of mortgagors)		0.178^{***}
		(0.055)
LTI	2.912^{***}	-12.854^{***}
	(0.019)	(2.939))
House price	-0.024^{***}	-0.032^{**}
	(0.000)	(0.013)
Age	-0.136^{***}	0.250^{**}
	(0.003)	(0.123)
Mortgage term	1.010	
	(0.004)	
Change in net worth		$0.023*$
		(0.014)
Loan value		0.221^{***}
		(0.081)
Income		-0.500^{**}
		(0.234)
GVA change		-0.113
		(0.091)
Constant	49.19^{***}	19.660^{***}
	(0.169)	(7.917)
F-test	13,522	na
Observations	5,727,160	47,261
Note:		*p<0.1; **p<0.05; ***p<0.01

Each location has a neighbourhood, which is defined by a circle with radius 10km. All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level.

their goods and services to local residents. Concretely, for locally tradable businesses, one would not expect employment choices to depend on local cash-flow shocks. Following best practice (e.g. Mian and Sufi (2014)), I therefore perform a placebo test on establishments that fall under the locally tradable industrial classification outlined in Appendix 7.3. Table 7 shows that there are no statistically significant cash-flow effects for locally tradable firms for neighbourhoods of different sizes. This is also true even when refraining from clustering the standard errors.

The classification of locally non-tradable and tradable firms is challenging and worth probing. Firms are asked to report the industrial code that best reflects their activities but this is often prone to error or simply difficult to do for a firm engaging in multiple business activities. In the first three columns of Table 8 I split my locally non-tradable establishments into three sectors.[47] Column 1 shows the cash-flow effect is strong for the *food and drink* sector: the employment effect coefficient is around 150% of δ in the main results. This is unsurprising because this sector is likely to be most associated with discretionary income. Restaurants and bars are also much more likely to smoothly adjust their labour inputs in response to demand shocks. Column 2 suggests that the cash-flow effect might be equally strong for the *vehicle* sector. One plausible hypothesis for this is that the purchase and (inessential) repair of cars can be delayed. Cash-flow shocks from lower mortgage payments might have brought forward consumption in this sector. Finally, the third column suggests there are very limited employment effects in the *retail* sector. The retail sector is dominated by large supermarkets, which are likely to face relatively income-inelastic demand, with limited options to adjust the number of staff in the face of a shock.

Small, independent firm directors often use residential collateral to support their business ventures and this may not be picked up in my controls (e.g., Adelino, Schoar, and Severino (2015) and Bahaj, Foulis, and Pinter (2017)). To rule out this supply-side channel, I restrict my sample to chain stores and large establishments in the fourth and fifth columns of Table 8. Chain stores are defined as belonging to enterprises that have more than one establishment. Large establishments are defined as those that employed more than 25 people in 2009.[48] Both columns show strong cash-flow effects, which suggests that it is the demand channel driving my results. If the supply channel was important then we would expect to see strong cash-flow effects for small establishments, yet there is little apparent cash-flow effect in the sixth column.[49]

[47] See table notes for the constituent industrial classifications of each sector.

[48] That is, nodes that contained all establishments with fewer than 25 people are dropped, regardless of total employment at the node.

[49] The definition of small establishments in the Table 8 is those employing between 11 and 25 people (so, large enough that small staffing changes are a relatively small fraction of overall employment). Results for establishments that are smaller still are negative and insignificant.

Table 7: Locally Tradable Firms

	Annual employment growth, pp		
	(1)	(2)	(3)
	5km radius	10km radius	15km radius
Cash-flow shock (% of income)	−0.44	−2.59	−3.91
	(2.940)	(3.34)	(3.62)
Controls	Yes	Yes	Yes
Specification	OLS	OLS	OLS
Observations	43,719	46,088	46,311
Note:			*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level.

Table 8: Firm Types

	Annual employment growth, pp					
	(1)	(2)	(3)	(4)	(5)	(6)
	Food and drink sector	Vehicle sector	Retail sector	Chains	Large establishments	Small establishments
Cash-flow shock (% of income)	6.92***	7.09**	1.58	5.82***	5.98***	0.21
	(2.05)	(3.46)	(2.02)	(1.92)	(1.66)	(3.72)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS	OLS	OLS
Observations	26,747	15,907	31,188	17,115	12,133	8,799

Note: *p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level. The *Food and Drink* sector is made up of Restaurants and food service; Beverage serving; and Catering. The *Vehicle* sector is made up of Sale of new motor vehicles; Sale of motor vehicle parts and accessories; Sale, maintenance and repair of motorcycles and related parts and accessories; and Retail sale of automotive fuel. The *Retail* sector is made up of (Retail sale of/in) Non-specialised stores; Fruit and vegetables; Medical and cosmetic goods; Specialised goods; Books, newspapers, recreation and stationery; and Repairs of personal and household goods. Chains are defined as enterprises with multiple establishments. Large establishments employ more than 25 people and small establishments employ between 11 and 25. Results are also insignificant for establishments employing fewer than 11 people.

5.3 Neighbourhood Definitions

I argue that one of the contributions of this paper is to move away from using large, contiguous administrative regions and construct neighbourhoods that better reflect where customers of locally non-tradable firms live. The calibration of these catchment areas is fraught with difficulty because there is a great deal of heterogeneity among locally non-tradable firms but also because local geography plays a subtle role. In Table 9 I show the main results are robust to the choice of neighbourhoods I use.

In the first two columns I vary the radius of the circular neighbourhood. The fact the coefficients do not change dramatically provides some reassurance that 10km is a reasonable compromise. Given that much of the effect comes from employment changes at large establishments, and these are likely to have larger customer catchment areas, it is small wonder the 15km radius neighbourhood is more statistically significant than the 5km radius. In the third and fourth columns I change the measure of my neighbourhoods from area to the number of people, which I have taken from postcode-level census data. Such a definition is likely to better reflect population density heterogeneity across the country but choosing an optimal size is still far from trivial. Once again, the coefficients are relatively stable and statistically significant, albeit around three quarters of the magnitude of the results using 10km radius neighbourhoods.

Finally, in last two columns I use official contiguous boundaries. In the fifth column I use local authority administrative regions, which is most analogous to the methodology of other studies (e.g. Mian and Sufi (2014)). Reduced precision in identification in these administrative regions leads to attenuation bias and a coefficient closer zero. If individual mortgage choice was purely determined by chance then these regions are large enough that we would expect there to be almost no variation in the aggregate proportion of variable-rate mortgages. In the sixth column I use Middle-layer Super Output Area (MSOA) boundaries. These are geographic areas designed to improve statistical reporting by grouping similar economic activity together.[50] But these suffer from the opposite problem in that they are very small and the catchment area of locally non-tradable firms almost certainly extends beyond their boundaries. The coefficient, roughly a quarter of the magnitude of those found in the main results, can be thought of in the same vein as that of the 5km radius neighbourhood result. The artificially low results for the two most common official boundaries used in the UK therefore motivates the new methodology used in the main results of this paper.

[50] These are the closest UK boundaries to Metropolitan Statistical Areas (MSAs) used in the US. Some local authorities look like MSAs to the extent that the major town or city is at the centre. Although Travel To Work Areas (TTWAs) exist in the UK, this concept does not aid my identification: where people work and where they shop are often different places.

Table 9: Neighbourhood Specifications

	Annual employment growth, pp					
	(1)	(2)	(3)	(4)	(5)	(6)
	5km radius	15km radius	100,000 people	150,000 people	Local authorities	MSOAs
Cash-flow shock (% of income)	2.41**	4.35***	3.25**	3.70***	1.65**	1.11***
	(1.21)	(1.60)	(1.35)	(1.39)	(0.70)	(0.30)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS	OLS	OLS
Observations	47,261	47,261	47,261	47,261	387	8,457

Note: *p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level.

5.4 Firm Dynamics

Figure 13 shows the distribution of employment changes between 2009 and 2010 for establishments that were present in both years. As a natural consequence of the distribution of employment, changes are strongly clustered around zero. Establishments accounting for about a quarter of employment did not change their net staffing position between 2009 and 2010, while those accounting for half of employment changed employment by fewer than 5 people in either direction.

Of the total fall in locally non-tradable employment, around half is accounted for by the net movement in staffing at this intensive margin and the other half was due to the net effect of firm creation and destruction. Figure 14 shows a decomposition of the 100,000 jobs that were lost in the locally non-tradable sector. It shows that the growth of small establishments positively contributed to employment but this was more than offset by large firms shrinking and small firms ceasing to exist. Firm turnover is surprisingly high in the UK, with locally non-tradable births and deaths making up around a third of the total number of locally non-tradable establishments in any given year.

Figure 13: Establishment-Level Employment Change (2009 - 2010)

(a) Linear Scale **(b)** Log Scale

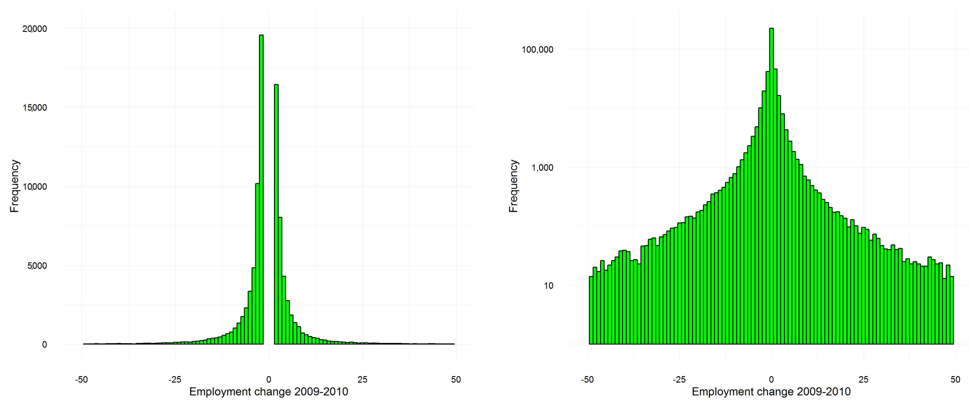

Source: ONS. Locally non-tradable firms only. Bars representing fewer than ten establishments have been excluded.

Together, this means that firm employment dynamics are more complicated than simple models might imply: gross flows far outweigh net flows, the modal employment growth of a establishment is zero and the population of establishments is continuously changing. Summing employment at each location helps to pin down the total effect on economy-wide employment growth but it captures both the extensive and intensive margin of labour adjustment. An alternative specification only sums up employment at each node for establishments that existed in both years - the intensive margin. This is consistent with studies such as Giroud and Mueller (2017) that use establishment-level regressions and log changes (since log changes at the extensive margin are undefined).

Figure 14: Total Employment Change (2009 - 2010)

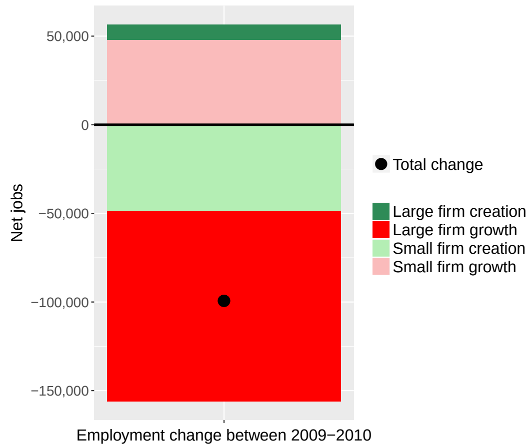

Source: ONS. This figure shows a decomposition of the total change in employment in the locally non-tradable sector between 2009 and 2010. Large firms are defined here as those with more than ten employees. Firm growth occurs when the establishment employs a non-zero number of people in both years. Firm creation is the net of establishment creation and establishment destruction.

Table 10 shows the results of the intensive-margin regressions. In the first column I find the cash-flow effect semi-elasticity is around 2.9pp and again significant at the 1-percent level. The second column shows the effect is only slightly weaker when ruling out the supply-side channel and only using chain establishments as before. Once again in the third column, the cash-flow effect is not picked up for non-chain establishments. Together with Figure 14, this suggests that we can attribute between half and two thirds of the cash-flow effect to firms making marginal changes to employment.

Table 10: Intensive Margin of Adjustment

	Annual employment growth, pp		
	(1)	(2)	(3)
	All firms	Multi-establishment	Single-establishment
Cash-flow shock (% of income)	2.90***	2.49**	0.706
	(0.84)	(1.00)	(0.87)
Controls	Yes	Yes	Yes
Specification	OLS	OLS	OLS
Observations	43,281	15,900	41,957
Note:			*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level.

6 Conclusion

I find that the monetary easing by the Bank of England in the autumn of 2008 had a significant and immediate cash-flow impact on people with a variable-rate mortgages. In areas where this cash-flow shock was especially large, it supported spending, including consumption for locally provided goods and services, thereby supporting employment in these sectors. Although monetary policy works through multiple channels with long and variable lags, I find the cash-flow channel through to locally non-tradable employment is sizable at the establishment level and leads to heterogeneity of policy impact across regions.

The analysis in this paper suggests that a 1pp accommodative monetary policy change was associated with a mortgagor cash-flow shock of around 0.75% of gross income, and around a 3.5pp increase in the annual employment growth for locally non-tradable establishments between 2009 and 2010. Locally non-tradable firms account for about a fifth of employment, so the total impact on employment following the 400bp cut in policy rates was around 2pp. This is a substantial estimated effect but applying cross-sectional identification to an aggregate effect requires some strong assumptions, many of which are laid out in Nakamura and Steinsson (2017).

To put the results in context, we can compare them to other estimates of the (overall) impact of monetary policy on employment using more standard techniques. Using an identification approach similar to Romer and Romer (2004), Cloyne and Hürtgen (2016) estimate that a 1pp transitory monetary policy shock leads to a peak decline in output of 0.6% and a 1pp fall in inflation. But they also show that the same shock leads to roughly a 10-20bp increase in unemployment at the 12-15 month horizon.[51] My results are larger, but the overall takeaway should be that local employment cash-flow effects are economically significant, even if the central estimates might be inappropriate for calibrating monetary policy decisions in general equilibrium.

The consequence of my estimated cash-flow effect is that monetary policy can lead to heterogeneous employment effects across space, as well as time. Regions that had a high proportion of people on large variable-rate mortgages, and employed a large fraction of their labour force in the locally non-tradable sector, benefitted the most from the first-round effects of accommodative monetary policy. To the extent there are significant differences in the mortgage markets across countries, this paper also sheds light on how the transmission of monetary policy might vary across the globe. But my work also has implications for the traction of monetary policy over time. Since policy rates reached their nadir in 2009, the average fixed-rate duration of new mortgages has increased from just over one year to close to four as people have tried their best to lock in low rates for an extended period of time. Although the duration of the stock is likely to adjust more slowly, this suggests the average household will experience smaller direct cash-flow

[51]The 68% confidence interval is between about 0.05 and 0.2pp; the 95% confidence interval is between just-below zero and around 0.3.

effects in the face of rising interest rates in the future.

References

Adams, R., H. Almeida, and D. Ferreira (2009): "Understanding the relationship between founder-CEOs and firm performance," *Journal of Empirical Finance*, 16(1), 136–150.

Adelino, M., A. Schoar, and F. Severino (2015): "House prices, collateral, and self-employment," *Journal of Financial Economics*, 117(2), 288–306.

Agarwal, S., I. Ben-David, and V. Yao (2017): "Systematic Mistakes in the Mortgage Market and Lack of Financial Sophistication," *Journal of Financial Economics*, 123(1), 42–58.

Agarwal, S., S. Chomsisengphet, C. Liu, and N. S. Souleles (2015): "Do Consumers Choose the Right Credit Contracts?," *The Review of Corporate Finance Studies*, 4(2), 239–257.

Agarwal, S., and B. Mazumder (2013): "Cognitive Abilities and Household Financial Decision Making," *American Economic Journal: Applied Economics*, 5(1), 193–207.

Aron, J., J. V. Duca, J. Muellbauer, K. Murata, and A. Murphy (2012): "Credit, Housing Collateral and Consumption: Evidence from the UK, Japan and the US," *Review of Income and Wealth*, 58(3), 397–423.

Aron, J., and J. Muellbauer (2016): "Modelling and Forecasting Mortgage Delinquency and Foreclosure in the UK," *Journal of Urban Economics*, 94, 32–53.

Auclert, A. (2017): "Monetary Policy and the Redistribution Channel," Working Paper 23451, National Bureau of Economic Research.

Bacon, P. M., and P. Moffatt (2012): "Mortgage Choice as a Natural Field Experiment on Choice under Risk," *Journal of Money, Credit and Banking*, 44(7), 1401–1426.

Badarinza, C., J. Y. Campbell, and T. Ramadorai (2017): "What Calls to ARMs? International Evidence on Interest Rates and the Choice of Adjustable-Rate Mortgages," *Management Science*.

Bahaj, S. A., A. Foulis, and G. Pinter (2017): "Home Values and Firm Behaviour," Discussion paper, Centre for Macroeconomics.

Bank of England (1999): "The Transmission Mechanism of Monetary Policy," *Bank of England Quarterly Bulletin*, 39(2), 161–170.

BERAJA, M., A. FUSTER, E. HURST, AND J. VAVRA (2017): "Regional Heterogeneity and Monetary Policy," Discussion paper, National Bureau of Economic Research.

BERNANKE, B. S. (2007): "The Financial Accelerator and the Credit Channel," Speech 296, Board of Governors of the Federal Reserve System.

BEST, M. C., J. CLOYNE, E. ILZETZKI, AND H. J. KLEVEN (2015): "Interest Rates, Debt and Intertemporal Allocation: Evidence from Notched Mortgage Contracts in the United Kingdom," Bank of England working papers 543, Bank of England.

BULLARD, J. (2014): "Income inequality and monetary policy: a framework with answers to three questions," Speech 235, Federal Reserve Bank of St. Louis.

BUNN, P., A. PUGH, AND C. YEATES (2018): "The Distributional Impact of Monetary Policy Easing in the UK between 2008 and 2014," Bank of England working papers 720, Bank of England.

BUNN, P., AND M. ROSTOM (2015): "Household Debt and Spending in the United Kingdom," Bank of England working papers 554, Bank of England.

BUNN, P., M. ROSTOM, S. DOMIT, N. WORROW, AND L. PISCITELLI (2013): "The Financial Position of British Households: Evidence from the 2013 NMG Consulting Survey," *Bank of England Quarterly Bulletin*, 53(4), 351–360.

BYRNE, D., R. KELLY, AND C. O'TOOLE (2017): "How Does Monetary Policy Pass-Through Affect Mortgage Default? Evidence from the Irish Mortgage Market," Research Technical Papers 04/RT/17, Central Bank of Ireland.

CALZA, A., T. MONACELLI, AND L. STRACCA (2013): "Housing Finance and Monetary Policy," *Journal of the European Economic Association*, 11(suppl_1), 101–122.

CAMPBELL, J. Y., AND J. F. COCCO (2003): "Household Risk Management and Optimal Mortgage Choice," *The Quarterly Journal of Economics*, 118(4), 1449–1494.

CAMPOS, C. (2012): "The Geographical Concentration of Industries," Regional economic analysis papers, Office of National Statistics.

CARLINO, G. A., AND R. H. DEFINA (1997): "The Differential Regional Effects of Monetary Policy: Evidence from the U.S. States," Working Papers 97-12, Federal Reserve Bank of Philadelphia.

CAVA, G. L., H. HUGHSON, AND G. KAPLAN (2016): "The Household Cash Flow Channel of Monetary Policy," RBA Research Discussion Papers rdp2016-12, Reserve Bank of Australia.

Chakraborty, C., M. Gimpelewicz, and A. Uluc (2017): "A Tiger by the Tail: Estimating the UK Mortgage Market Vulnerabilities from Loan-Level Data," Bank of England working papers 703, Bank of England.

Cloyne, J., C. Ferreira, and P. Surico (2016): "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism," Discussion paper.

Cloyne, J., K. Huber, E. Ilzetzki, and H. Kleven (2017): "The Effect of House Prices on Household Borrowing: A New Approach," Working Paper 23861, National Bureau of Economic Research.

Cloyne, J., and P. Hürtgen (2016): "The Macroeconomic Effects of Monetary Policy: A New Measure for the United Kingdom," *American Economic Journal: Macroeconomics*, 8(4), 75–102.

Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2017): "Innocent Bystanders? Monetary policy and Inequality," *Journal of Monetary Economics*, 88, 70 – 89.

Corsetti, G., J. B. Duarte, and S. Mann (2018): "One Money, Many Markets - A Factor Model Approach to Monetary Policy in the Euro Area with High-Frequency Identification," Cambridge working papers in economics, Faculty of Economics, University of Cambridge.

Davis, S. J., J. Haltiwanger, R. Jarmin, and J. Miranda (2007): "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," in *NBER Macroeconomics Annual 2006, Volume 21*, NBER Chapters, pp. 107–180. National Bureau of Economic Research, Inc.

Department for Transport (2016): "National Travel Survey: England 2015," Discussion paper, Statistical Release.

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017): "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging," *American Economic Review*, 107(11), 3550–88.

Dupor, B., and R. Guerrero (2017): "Local and Aggregate Fiscal Policy Multipliers," *Journal of Monetary Economics*, 92, 16 – 30.

Flodén, M., M. Kilström, J. Sigurdsson, and R. Vestman (2017): "Household Debt and Monetary Policy: Revealing the Cash-Flow Channel," Discussion paper.

Fornero, E., C. Monticone, and S. Trucchi (2011): "The Effect of Financial Literacy on Mortgage Choices," .

Fuster, A., and P. S. Willen (2013): "Payment Size, Negative Equity, and Mortgage Default," Working Paper 19345, National Bureau of Economic Research.

Giroud, X., and H. M. Mueller (2017): "Firms' Internal Networks and Local Economic Shocks," Discussion paper, National Bureau of Economic Research.

Guren, A. M., A. Krishnamurthy, and T. J. McQuade (2018): "Mortgage Design in an Equilibrium Model of the Housing Market," Discussion paper, National Bureau of Economic Research.

Hellebrandt, T., S. Pezzini, J. Saleheen, and R. Williams (2009): "The Financial Position of British Households: Evidence from the 2009 NMG Survey," *Bank of England Quarterly Bulletin*, 49(4), 274–285.

Jackman, R., and J. Sutton (1982): "Imperfect Capital Markets and the Monetarist Black Box: Liquidity Constraints, Inflation and the Asymmetric Effects of Interest Rate Policy," *The Economic Journal*, 92(365), 108–128.

Kerr, J., L. Frank, J. F. Sallis, B. Saelens, K. Glanz, and J. Chapman (2012): "Predictors of Trips to Food Destinations," *International Journal of Behavioral Nutrition and Physical Activity*, 9(1), 58.

Luck, S., and T. Zimmermann (2017): "Employment Effects of Unconventional Monetary Policy: Evidence from QE," .

Mian, A., and A. Sufi (2014): "What Explains the 2007 2009 Drop in Employment?," *Econometrica*, 82(6), 2197–2223.

Miles, D. (2005): "Incentives Information and Efficiency in the UK Mortgage Market," *The Economic Journal*, 115, C82–C98.

Nakamura, E., and J. Steinsson (2014): "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 104(3), 753–92.

——— (2017): "Identification in Macroeconomics," Working Paper 23968, National Bureau of Economic Research.

Ozkan, S., K. Mitman, F. Karahan, and A. Hedlund (2016): "Monetary Policy, Heterogeneity and the Housing Channel," 2016 Meeting Papers 663, Society for Economic Dynamics.

Piskorski, T., and A. Seru (2018): "Mortgage Market Design: Lessons from the Great Recession," Discussion paper, Brookings papers on economic activity.

Piskorski, T., and A. Tchistyi (2010): "Optimal Mortgage Design," *The Review of Financial Studies*, 23(8), 3098–3140.

Romer, C. D., and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94(4), 1055–1084.

Rubio, M. (2011): "Fixed- and Variable-Rate Mortgages, Business Cycles, and Monetary Policy," *Journal of Money, Credit and Banking*, 43(4), 657–688.

Selezneva, V., M. Schneider, and M. Doepke (2015): "Distributional Effects of Monetary Policy," 2015 Meeting Papers 1099, Society for Economic Dynamics.

Violante, G., G. Kaplan, and J. Weidner (2014): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 45(1 (Spring)), 77–153.

7 Appendix

7.1 Mortgage stock construction

The PSD provides information on mortgage issuance flows but I need to estimate characteristics of the stock of mortgages in the autumn of 2008.[52] I therefore follow the steps outlined in Chakraborty, Gimpelewicz, and Uluc (2017), which identify the mortgages from the flow that still exist in late 2008 and go on to model how the interest rates and quantities pertaining to the loans evolved. Estimating the stock is more complicated than summing the flows for a few reasons. First, I need to remove mortgages that have been paid off in full. It is relatively easy to remove mortgages with a very short maturity but it is more challenging to find the small number of mortgages that are paid down through a windfall gain (e.g. inheritance). I also need to remove loans that have defaulted, which is again difficult to do with full precision but a good proxy is to strip out loans at the very top of the LTI distribution given the evidence that these are most vulnerable to default (e.g., Aron and Muellbauer (2016)).[53] Removing redundant loans after a remortgage is relatively straightforward and requires matching birth dates and postcodes before keeping all but the most recent entry before July 2008.

Finally, I need to remove loans that were paid off in the process of moving house as many mortgages in the UK are not portable between properties. The majority of redundant mortgages are found by looking for three-part tuples of transactions. The first is the entry of the home-mover when they bought their original property (person a, property x, time $t-1$). The second is for the home-mover at the point they bought their new property (person a, property y, time t). The third is the entry corresponding to the person who moved into the home-mover's old house (person b, property x, time t). The key is therefore to match the birth date for the first two parts and the timing (to the nearest few days) for the last two parts.[54] After this process is complete, I am left with around 5.8m mortgages. Figure 15 provides an estimate of the mortgages that might be missing from the sample.

Figure 15 shows an estimate of the number of mortgages that might be missing from the stock used in the main analysis. Regulatory data captures mortgage flows after April 2005 so the vast majority of missing mortgages were issued before this time. Reliable regulatory data for the stock of mortgages is available from 2015Q4 and this suggests there were around 500,000 mortgages issued before 2005Q2 (and still existed in 2015Q4). Assuming the 2015Q4 stock of around 6.7m mortgages was relatively constant over time, that means there were likely a further

[52] The PSD only provides information on the stock of mortgages after 2015.

[53] Between 2005 and 2008 only 100,000 properties were taken into court possession in England and Wales; that includes properties purchased with buy-to-let mortgages. See https://www.bsa.org.uk/statistics/mortgages-housing.

[54] Some of these chains are incomplete so the remaining redundant mortgages are found using a bi-partite graph-matching algorithm. For full details see Chakraborty, Gimpelewicz, and Uluc (2017).

Figure 15: Missing Mortgages

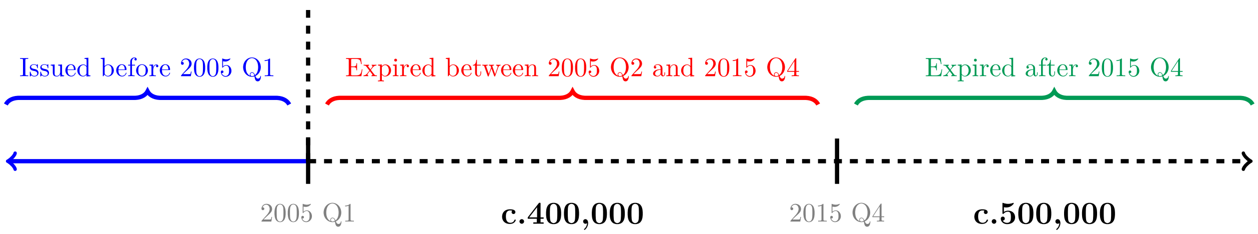

400,000 mortgages issued before 2005Q2 that expired between 2008Q3 and 2015Q4. If anything, the number of mortgagors has actually increased over time so these estimates for the numbers of missing mortgages are likely to be conservatively high.

7.2 Imputing Missing Variables

In addition to the variables with complete coverage there are also a few additional categories that are occasionally left incomplete due to field changes or decisions by lenders not to report. The two most important fields that are left blank are the initial interest rate of the loan and the contractual maturity of the interest rate. In the UK there are very few mortgages with fixed interest rates spanning the duration of the mortgage term. Instead, the most common types of mortgage are either fixed or linked to official policy interest rates for a relatively short period of time (usually between two and five years). Absence a refinance at this point, they revert to the lender's SVR, which is linked to official policy interest rates and was often unattractively high before 2007.[55]

To overcome the lack of information on the initial interest rate, I follow Best, Cloyne, Ilzetzki, and Kleven (2015) and exploit the fact that the UK mortgage market is very competitive, which means that mortgage interest rates are predictable to a high degree of accuracy when one controls for the other mortgage and borrower characteristics available in the PSD. It is useful to run regressions separately by year, y and mortgage *type* (first time buyer, home mover and remortgagor) of the form:

$$\begin{aligned} r_i^{y,type} &= LTV_i + \beta_1 \times deal_i + \beta_2 \times ratetype_i + \beta_3 \times repayment_i \\ &+ \beta_4 \times income_i\{single_i\} + \beta_5 \times income_i\{joint_i\} + \varepsilon_i \end{aligned} \quad (8)$$

This is a modified version of Equation 1 in Best, Cloyne, Ilzetzki, and Kleven (2015). The dependent variable is the initial mortgage interest rate for individual i. On the right hand side

[55] As shown below, the SVR was often an attractive option during 2009 when spreads on remortgages increased at the same time that policy interest rates fell.

LTV_i is a dummy for the threshold LTV ratio of the mortgage, *deal* is a vector of dummies for the contractual term (two, three, five or ten years), *ratetype* is a dummy for how the interest rate evolves during the initial contract term (fixed or variable) and *repayment* is a dummy for whether the mortgage amortises or is interest-only. The last two independent variables interact the reported gross income with the *income* basis reported (single or joint). My empirical approach outlined below relies on the dramatic change in interest rates at the end of 2008, which means the estimated interest rate levels have a relatively wide tolerance for their accuracy.[56] Nevertheless, the above specification does a good job at forecasting an approximate initial interest rate when the information is available; 95% confidence intervals are almost invisible in Figure 16.

The regressions for first time buyers and remortgagors are reported in Tables 11 and 12. It is straightforward from these to compute conditional LTV-interest rate curves by taking the average values of the covariates to give an indication of how mortgage rates varied by leverage before the Crisis, as shown in Figure 16. The main specification uses these imputed interest rates when they are missing, though the majority of mortgages with missing interest rate information are variable-rate contracts where the interest rate change (rather than the starting level) is relatively more important for the cash-flow shock. A robustness test run using only mortgages with complete entries across all data fields leads to similar findings.

A second key variable that is sometimes missing from my data is the contractual maturity. This is important because, although all lenders report whether the initial deal is fixed of variable, the contractual term determines when people move onto the SVR.[57] For example, if the contractual term is unknown, we do not know if a mortgagor with a fixed-rate mortgage issued in 2006 is part-way through the contractual period in mid 2009 or if the interest rate had already reverted to the SVR. In the latter case, the mortgagor would have benefitted slightly from the fall in interest rates during the 2008Q3-2009Q4 window of interest. In the main specification I assume that mortgagors are on a two-year contractual term if this information is left blank as this was the modal contract on offer in 2008.[58] The vast proportion of all other mortgages at this time were three-year mortgages. I have performed a robustness check assuming all the missing contractual maturities were three instead of two years and find quantitatively similar estimates for β and δ as in the main regressions.

[56]This estimation procedure helps with modelling the amortisation of the mortgage and the effect of the shock, though estimates are quantitatively similar if an approximate initial rate of Bank Rate plus 150bp is used for all mortgages.

[57]Around four million mortgages have no associated contractual maturity in my estimated stock. This data field was filled out more frequently in later quarters as reporting oversight became stricter.

[58]Using information available, between a half and two thirds of mortgages were two-year fixed rate contracts before 2008. It is also intuitive that lenders would be more likely to leave information blank if it was the most common product they were offering.

Figure 16: Pre-Crisis LTV Curves

(a) First Time Buyers **(b)** Remortgagors

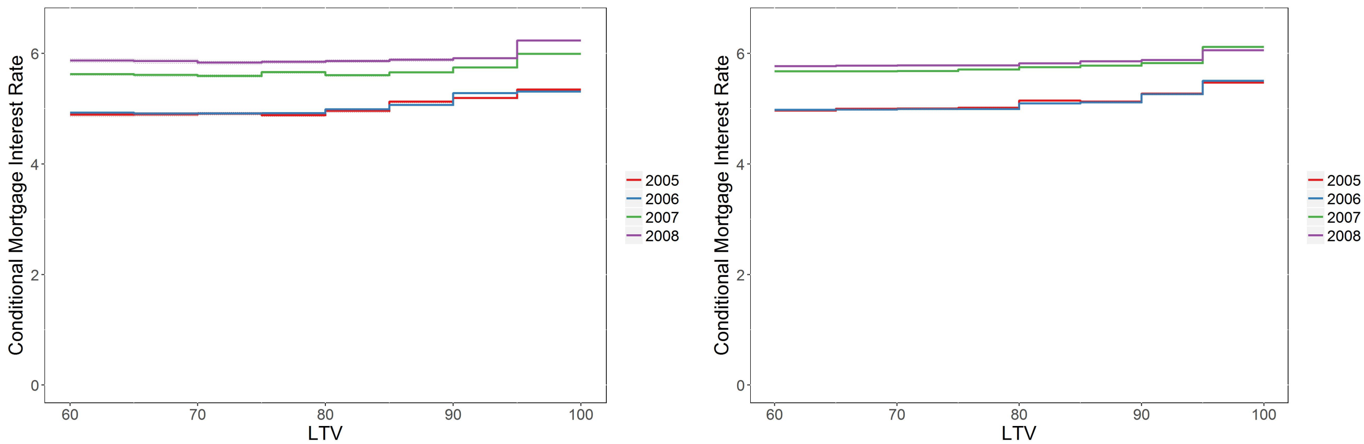

Theses charts use coefficients in Tables 11 and 12 to give an indication of mortgage interest rates offered, conditional on mean values for the covariates.

7.3 Firm Classifications

In 2010 there were just under 28.5m people employed in the UK. Around 3m were employed by local authorities and 2.4m employed by central government. Of the remaining 23m private-sector employees, around 4.7m were employed in what I define as locally non-tradable firms and 2.2m in locally tradable firms. This appendix gives some more details about those definitions.

This mechanical mapping from the US industrial categories leads to similar results to the spatial analysis carried out for the UK by Campos (2012). Of the thirty least geographically concentrated industries in the UK, only 11 relate to non-public and non-construction activities. The non-tradable definition used in the rest of this analysis captures all but two of these industries: wholesale activities and transport systems are excluded because they are unlikely to effectively capture local demand effects convincingly. Of the thirty most geographically concentrated industries in the UK, about half are captured in the main mapping from the existing literature; most are firms involved in some form of manufacturing. The other half contain industries associated with finance, transport, holiday recreation and professional services - many of which would plausibly fall under the intuitive definition of tradables in that the firms do not garner the majority of their sales from locally resident customers.

According to the primary classification system used in this study, locally non-tradable industry employment makes up 21% of the UK aggregate private sector employment, which is very similar to the share in the US. Of the 14 locally non-tradable industry groups, the largest employers are retailers of groceries, restaurants and other general merchandise stores. Unsurprisingly, these outlets tend to be concentrated in the most urban regions of the UK, in city centres

Table 11: First Time Buyer Interest Rates

	Dependent variable:			
	Conditional interest rate			
	2005	2006	2007	2008
	(1)	(2)	(3)	(4)
LTV 55-60%	-0.093^{***}	-0.120^{***}	-0.222^{***}	-0.282^{***}
	(0.015)	(0.010)	(0.011)	(0.017)
LTV 60-65%	-0.085^{***}	-0.133^{***}	-0.236^{***}	-0.292^{***}
	(0.014)	(0.010)	(0.011)	(0.017)
LTV 65-70%	-0.070^{***}	-0.132^{***}	-0.252^{***}	-0.320^{***}
	(0.012)	(0.009)	(0.010)	(0.016)
LTV 70-75%	-0.099^{***}	-0.131^{***}	-0.186^{***}	-0.306^{***}
	(0.011)	(0.008)	(0.009)	(0.013)
LTV 75-80%	-0.021^{*}	-0.060^{***}	-0.241^{***}	-0.290^{***}
	(0.011)	(0.008)	(0.008)	(0.013)
LTV 80-85%	0.146^{***}	0.020^{***}	-0.187^{***}	-0.267^{***}
	(0.010)	(0.007)	(0.007)	(0.011)
LTV 85-90%	0.210^{***}	0.235^{***}	-0.100^{***}	-0.238^{***}
	(0.007)	(0.005)	(0.005)	(0.009)
LTV 90-95%	0.362^{***}	0.263^{***}	0.145^{***}	0.081^{***}
	(0.007)	(0.005)	(0.005)	(0.008)
1y deal	0.016	-0.037^{***}	0.070^{***}	-0.168^{***}
2y deal	0.164^{***}	0.228^{***}	0.139^{***}	0.083^{***}
3y deal	0.454^{***}	0.471^{***}	0.170^{***}	0.107^{***}
4y deal	0.896^{***}	0.557^{***}	0.008	0.039
5y+ deal	0.525^{***}	0.660^{***}	0.402^{***}	0.033^{***}
Fixed rate	0.172^{***}	0.183^{***}	0.084^{**}	0.484^{***}
Tracker	0.661^{***}	0.783^{***}	0.610^{***}	0.914^{***}
SVR	0.024	0.136^{***}	0.164^{***}	0.229^{***}
Interest only	-0.566^{***}	-0.231^{***}	0.002	0.123^{***}
Single income	-0.00003^{***}	-0.001^{***}	-0.0004^{***}	-0.0001^{***}
Joint income	-0.001^{***}	-0.001^{***}	-0.001^{***}	-0.001^{***}
Constant	4.726^{***}	4.753^{***}	5.677^{***}	5.695^{***}
Observations	118,105	176,087	156,697	47,321
R^2	0.213	0.258	0.131	0.152
Residual Std. Error	0.638 (df = 118085)	0.554 (df = 176067)	0.576 (df = 156677)	0.548 (df = 47301)
F Statistic	$1{,}682.215^{***}$ (df = 19; 118085)	$3{,}219.678^{***}$ (df = 19; 176067)	$1{,}246.575^{***}$ (df = 19; 156677)	444.995^{***} (df = 19; 47301)
Note:				*p<0.1; **p<0.05; ***p<0.01

The dependent variable is the initial interest rate offered on the mortgage. The main explanatory dummy variables are the loan-to-value bucket, the term of the deal (1,2,3,4,5+ years), the interest rate type (tracker, fixed or SVR), the amortisation status and the income of the mortgagor (an interaction between gross income and whether it was a single or joint mortgage). The dummies show the marginal effect on the interest rate over and above a baseline amortising mortgage with an LTV of less than 55%, a contractual term of less than one year and a capped interest rate.

Table 12: Remortgagor Interest Rates

	Dependent variable:			
	Conditional interest rate			
	2005	2006	2007	2008
	(1)	(2)	(3)	(4)
LTV 55-60%	−0.015***	−0.047***	−0.109***	−0.070***
	(0.004)	(0.003)	(0.004)	(0.004)
LTV 60-65%	0.015***	−0.043***	−0.108***	−0.064***
	(0.004)	(0.003)	(0.004)	(0.004)
LTV 65-70%	0.022***	−0.033***	−0.101***	−0.060***
	(0.004)	(0.003)	(0.004)	(0.004)
LTV 70-75%	0.034***	−0.032***	−0.077***	−0.059***
	(0.004)	(0.003)	(0.004)	(0.004)
LTV 75-80%	0.163***	0.072***	−0.033***	−0.022***
	(0.005)	(0.003)	(0.004)	(0.004)
LTV 80-85%	0.149***	0.090***	−0.005	0.016***
	(0.004)	(0.003)	(0.004)	(0.004)
LTV 85-90%	0.291***	0.236***	0.041***	0.040***
	(0.005)	(0.003)	(0.004)	(0.004)
LTV 90-95%	0.490***	0.477***	0.335***	0.214***
	(0.006)	(0.004)	(0.005)	(0.006)
1y deal	0.355***	0.001	−0.231***	−0.104***
2y deal	0.023***	−0.025***	0.010***	0.104***
3y deal	0.423***	0.345***	0.243***	0.072***
4y deal	0.286***	0.249***	−0.021*	0.255***
5y+ deal	0.169***	0.157***	0.119***	0.102***
Fixed rate	0.206***	0.055***	−0.153***	−0.044***
Tracker	0.975***	0.642***	0.424***	0.571***
SVR	0.232***	0.020***	−0.005	−0.098***
Interest only	−0.239***	−0.142***	0.0003	0.044***
Single income	−0.0005***	−0.0005***	−0.0002***	−0.0002***
Joint income	−0.0001***	−0.001***	−0.0003***	−0.001***
Constant	4.691***	4.952***	5.862***	5.872***
Observations	323,448	491,814	452,083	219,947
R^2	0.155	0.137	0.054	0.076
Residual Std. Error	0.582 (df = 323428)	0.484 (df = 491794)	0.626 (df = 452063)	0.457 (df = 219927)
F Statistic	3,132.035*** (df = 19; 323428)	4,117.412*** (df = 19; 491794)	1,345.967*** (df = 19; 452063)	957.108*** (df = 19; 219927)
Note:				*p<0.1; **p<0.05; ***p<0.01

The dependent variable is the initial interest rate offered on the mortgage. The main explanatory dummy variables are the loan-to-value bucket, the term of the deal (1,2,3,4,5+ years), the interest rate type (tracker, fixed or SVR), the amortisation status and the income of the mortgagor (an interaction between gross income and whether it was a single or joint mortgage). The dummies show the marginal effect on the interest rate over and above a baseline amortising mortgage with an LTV of less than 55%, a contractual term of less than one year and a capped interest rate.

such as London, Birmingham and Leeds but also in the suburban commuter zones and tourist destinations.

The first two columns in the following tables present the 25 locally non-tradable and 82 tradable classifications using the NAICS-12 system used in Mian and Sufi (2014). The third and fourth columns provide the closest matches to the SIC-03 system. The final column provides the share of employment as a proportion of total locally non-tradable and tradable total employment, respectively.

The vintage of my data makes it is necessary to create a mapping between NAICS-12 and SIC-07, then back to SIC-03. Some NAICS-12 industrial classes therefore correspond to multiple SIC-03 codes. In the penultimate column of the tables I have listed the closest code matches.

7.3.1 Locally Non-Tradable Firms

Table 13: Locally Non-Tradable Industry Definitions

NAICS-12 description	NAICS-12 code	SIC-03 description	SIC-03 code	Employment share, %
Automobile Dealers	4411	Sale of new motor vehicles	501	4.7
Other Motor Vehicle Dealers	4412	Sale of motor vehicle parts and accessories	503	2.0
Automotive Parts, Accessories, and Tire Stores	4413	Sale, maintenance and repair of motorcycles and related parts and accessories	504	0.2
Furniture Stores	4421	Retail sale of automotive fuel	505	0.9
Home Furnishings Stores	4422	Retail sale in non-specialised stores	521	30.3
Electronics and Appliance Stores	4431	Retail sale of fruit and vegetables	522	3.6
Grocery Stores	4451	Retail of Medical and cosmetic	523	2.2
Specialty Food Stores	4452	Retail of specialised goods	524	25.1
Beer, Wine, and Liquor Stores	4453	Retail sale of books, newspapers, recreation and stationery	525	0.5
Health and Personal Care Stores	4461	Repair of personal and household goods	527	0.5
Gasoline Stations	4471	Restaurants and food service	553	14.7
Clothing Stores	4481	Beverage serving	554	10.7
Shoe Stores	4482	Catering	555	4.6
Jewelry, Luggage, and Leather Goods Stores	4483			
Sporting Goods, Hobby, and Musical Instrument Stores	4511			
Book Stores and News Dealers	4512			
Department Stores	4521			
Other General Merchandise Stores	4529			
Florists	4531			
Office Supplies, Stationery, and Gift Stores	4532			
Used Merchandise Stores	4533			
Other Miscellaneous Store Retailers	4539			
Restaurants and Other Eating Places	7225			
Special Food Services	7223			
Drinking Places (Alcoholic Beverages)	7224			

7.3.2 Locally Tradable Firms

Table 14: Locally Tradable Industry Definitions

NAICS-12 description	NAICS-12 code	SIC-03 description	SIC-03 code	Employment share, %
Forest nurseries and gathering of forest products	1132	Growing of fruit, nuts, beverage and spice crops	11	5.9
Fishing	1141	Fishing	50	0.4
Oil and gas extraction	2111	Deep coal mines and manufacture of solid fuel	101	0.3
Coal mining	2121	Mining and agglomeration of lignite	102	0.0
Metal ore mining	2122	Extraction of crude petroleum and natural gas	111	0.6
Nonmetallic mineral mining and quarrying	3111	Mining of uranium and thorium ores	120	0.0
Animal food manufacturing	3112	Mining of iron ores	131	0.2
Grain and oilseed milling	3113	Quarrying of ornamental and building stone	141	0.2
Sugar and confectionery product manufacturing	3114	Mining of chemical and fertilizer minerals	143	0.0
Fruit and vegetable preserving and specialty food manufacturing	3115	Slaughtering of animals other than poultry and rabbits	151	4.2
Dairy product manufacturing	3115	Freezing of fish	152	0.7
Animal slaughtering and processing	3116	Processing and preserving of potatoes	153	2.1
Seafood product preparation and packaging	3117	Liquid milk and cream production	155	1.0
Bakeries and tortilla manufacturing	3118	Grain milling	156	0.6
Other food manufacturing	3119	Manufacture of bread; manufacture of baked; manufacture of sugar	158	7.4
Beverage manufacturing	3121	Manufacture of distilled potable alcoholic beverages	159	1.7
Tobacco manufacturing	3122	Manufacture of tobacco products	160	0.3
Fiber, yarn, and thread mills	3131	Preparation and spinning of cotton-type fibres	171	0.2
Fabric mills	3132	Manufacture of knitted and crocheted fabrics	176	0.1
Textile and fabric finishing and fabric coating mills	3133	Manufacture of leather clothes	181	0.0
Textile furnishings mills	3141	Dressing and dyeing of fur; manufacture of fur articles	183	0.0
Other textile product mills	3149	Manufacture of footwear	193	0.2
Apparel knitting mills	3151	Manufacture of panels and boards	202	0.2
Cut and sew apparel manufacturing	3152	Manufacture of pulp	211	0.5
Apparel accessories and other apparel manufacturing	3159	Printing of newspapers	222	5.8
Leather and hide tanning and finishing	3161	Manufacture of coke oven products	231	0.0
Footwear manufacturing	3162	Manufacture of industrial gases	241	1.7
Other leather and allied product manufacturing	3169	Manufacture of pesticides and other agro-chemical products	242	0.1
Pulp, paper, and paperboard mills	3221	Manufacture of paints, varnishes and similar coatings	243	0.8
Converted paper product manufacturing	3222	Manufacture of basic pharmaceutical products	244	2.6
Printing and related support activities	3231	Manufacture of other chemical products	246	1.2
Petroleum and coal products manufacturing	3241	Manufacture of rubber tyres and tubes	251	1.1
Basic chemical manufacturing	3251	Manufacture of plastic plates, sheets, tubes and profiles	252	6.3

Resins and synthetic fibers and filaments manufacturing	3252	Manufacture of flat glass	261	1.0
Pesticide, fertilizer, and other agricultural chemical manufacturing	3253	Production of abrasive products	268	0.3
Pharmaceutical and medicine manufacturing	3254	Manufacture of basic iron and steel and of ferro-alloys	271	1.3
Paint, coating, and adhesive manufacturing	3255	Manufacture of steel tubes	272	0.4
Soap, cleaning compound, and toilet preparation manufacturing	3256	Cold drawing	273	0.1
Other chemical product and preparation manufacturing	3259	Manufacture of metal structures and parts of structures	281	3.6
Plastics product manufacturing	3261	Manufacture of central heating radiators and boilers	282	0.5
Rubber product manufacturing	3262	Manufacture of steel drums and similar containers	287	2.3
Clay product and refractory manufacturing	3271	Manufacture of non-vehicle engines and turbines	291	3.7
Glass and glass product manufacturing	3272	Manufacture of furnaces and furnace burners	292	4.2
Other nonmetallic mineral product manufacturing	3259	Manufacture of other machine tools	294	0.6
Iron and steel mills and ferroalloy manufacturing	3311	Manufacture of weapons and ammunition	296	0.8
Alumina and aluminum production and processing	3313	Manufacture of electric domestic appliances	297	0.7
Nonferrous metal (except aluminum) production and processing	3314	Manufacture of computers	300	1.0
Foundries	3315	Manufacture of electric motors, generators and transformers	311	1.1
Cutlery and handtool manufacturing	3322	Manufacture of insulated wire and cable	313	0.4
Boiler, tank, and shipping container manufacturing	3324	Manufacture of accumulators, primary cells and primary batteries	314	0.1
Hardware manufacturing	3325	Manufacture of lighting equipment and electric lamps	315	0.6
Spring and wire product manufacturing	3326	Manufacture of other electrical equipment	316	1.1
Machine shops and screw, nut, and bolt manufacturing	3327	Manufacture of audio and visual equipment	323	0.5
Other fabricated metal product manufacturing	3329	Manufacture of medical and surgical equipment and orthopaedic appliances	331	1.7
Agriculture, construction, and mining machinery manufacturing	3331	Manufacture of electronic instruments	332	2.6
Industrial machinery manufacturing	3332	Manufacture of motor vehicles	341	3.0
Commercial and service industry machinery manufacturing	3333	Manufacture of bodies (coachwork) for motor vehicles	342	1.0
Ventilation, heating and commercial refrigeration equipment manufacturing	3334	Building and repairing of ships	351	1.6
Metalworking machinery manufacturing	3335	Manufacture of other transport equipment	355	0.1
Engine, turbine, and power transmission equipment manufacturing	3336	Striking of coins	362	0.3
Other general purpose machinery manufacturing	3339	Manufacture of musical instruments	363	0.1
Computer and peripheral equipment manufacturing	3341	Manufacture of sports goods	364	0.2
Communications equipment manufacturing	3342	Manufacture of professional and arcade games and toys	365	0.2
Audio and video equipment manufacturing	3343	Manufacture of brooms and brushes	366	1.2
Semiconductor and other electronic component manufacturing	3344	Steam and hot water supply	403	0.0
Navigational, measuring, electromedical, and control instruments manufacturing	3345	Publishing of software	722	17.6

Manufacturing and reproducing magnetic and optical media	3346			
Electric lighting equipment manufacturing	3351			
Household appliance manufacturing	3352			
Electrical equipment manufacturing	3353			
Other electrical equipment and component manufacturing	3359			
Motor vehicle manufacturing	3361			
Motor vehicle body and trailer manufacturing	3362			
Motor vehicle parts manufacturing	3363			
Aerospace product and parts manufacturing	3364			
Railroad rolling stock manufacturing	3365			
Ship and boat building	3366			
Other transportation equipment manufacturing	3369			
Office furniture (including fixtures) manufacturing	3372			
Medical equipment and supplies manufacturing	3391			
Other miscellaneous manufacturing	3399			
Software publishers	5112			

7.4 Who switches mortgage type?

Table 1 showed the choice between taking out a fixed or variable-rate mortgage was relatively evenly balanced when examining the population of all mortgagors who refinanced once between 2005 and 2008. There was little evidence that, for the population as a whole, people sorted themselves into one mortgage type over another. In this appendix I explore this in more detail and show that the case against detrimental selection bias is even weaker.

The first two major columns of Table 15 split the population of one-time remortgagors into upper and lower quartiles according to their age. It is true that younger mortgagors are more likely to take out a fixed-rate mortgage. But, as before, the proportion of young mortgagors that switch mortgage type upon refinancing is reassuringly high. Conditional on no intertemporal relationship, we would expect to see around 41.6% of young mortgagors opting for the opposite contract upon refinance. In fact, we observe 34.8% flipping. Among both old and young mortgagors, the proportion of people staying on the original type of mortgage is slightly higher than expected. But the key is the bottom-left to top-right diagonals in the bottom panel are reasonably close to their expected values under the null hypothesis. The third and fourth major columns repeat the analysis for the LTI (which is a proxy for age) and a similar story holds.

The first two major columns of Table 16 split the population into their employment status. Self-employed mortgagors are slightly more likely to take out a fixed-rate mortgage. But for both the employed and self-employed, the bottom-left to top-right diagonals in the bottom panel add up to above 35%. The final concern I address is to rule out regional selection. In the final two major columns of Table 16 I take the regional extremes in terms of proportion of people opting for fixed-rate mortgages across time. Lower incomes in Scotland drive an overall preference for fixed-rate mortgages and the Northern Irish mortgage market idiosyncratically has more variable-rate mortgages compared to the rest of the UK. The proportion of mortgagors switching in both cases is within 10pp of the expected proportion under the null hypothesis. Together, this evidence strengthens the case made before that mortgage choice is driven by a number of factors (including the slope of the yield curve) and mortgage type does not simply reflect the mortgagor's inherent characteristics.

Table 15: Transitions: Age and LTI

Expected | no intertemporal relationship

Percent of remortgagors

	Second mortgage							
	Age below 31		Age above 44		LTI above 3.48		LTI below 2.25	
First mortgage	Variable	Fixed	Variable	Fixed	Variable	Fixed	Variable	Fixed
Variable	8.7	20.8	18.9	24.6	11.4	22.4	18.5	24.5
Fixed	20.8	49.8	24.6	32.0	22.4	43.8	24.5	32.5
Prob(Fix)	71%		57%		66%		57%	

Observed

Percent of remortgagors

	Second mortgage							
	Age below 31		Age above 44		LTI above 3.48		LTI below 2.25	
First mortgage	Variable	Fixed	Variable	Fixed	Variable	Fixed	Variable	Fixed
Variable	12.1	14.8	24.2	18.3	15.5	15.2	24.5	17.3
Fixed	20.0	53.2	20.2	37.3	21.5	47.9	19.7	38.5
Total mortgagors	207,547		191,671		192,632		188,753	

Age and LTI categories represent upper and lower quartiles as of the first mortgage. If the probability of choosing a fixed-rate contract was independent across time we would expect to see the distribution in the top panel. The observed distribution is shown in the bottom panel.

Table 16: Transitions: Regions and Employment Status

Expected | no intertemporal relationship

Percent of remortgagors

	Second mortgage							
	Self-employed		Employed		Scotland		Northern Ireland	
First mortgage	Variable	Fixed	Variable	Fixed	Variable	Fixed	Variable	Fixed
Variable	18.6	24.5	12.6	22.9	11.2	22.3	31.2	24.7
Fixed	24.5	32.4	22.9	41.5	22.3	44.3	24.7	19.5
Prob(Fix)	64%		57%		67%		44%	

Observed

Percent of remortgagors

	Second mortgage							
	Self-employed		Employed		Scotland		Northern Ireland	
First mortgage	Variable	Fixed	Variable	Fixed	Variable	Fixed	Variable	Fixed
Variable	22.7	19.1	17.5	16.3	16.1	16.8	34.1	25.2
Fixed	21.6	36.5	19.8	46.4	17.9	49.2	18.5	22.3
Total mortgagors	132,844		614,306		57,111		21,218	

Scotland is the region with the highest proportion of fixed-rate mortgages and Northern Ireland the lowest. If the probability of choosing a fixed-rate contract was independent across time we would expect to see the distribution in the top panel. The observed distribution is shown in the bottom panel.

7.5 Main Regression Table

Table 17: Main Specification

	Annual employment growth, pp		
	(1)	(2)	(3)
Variable rate (% of mortgagors)	0.188***		
	(0.059)		
Cash-flow shock (% of income)		4.646***	
		(1.437)	
Cash-flow shock (£1,000s)			6.814***
			(2.546)
LTI	−12.960*	−17.684**	−7.105
	(7.320)	(7.880)	(6.568)
House price	−0.034	−0.027	−0.056*
	(0.024)	(0.022)	(0.032)
Change in net worth	0.025	0.023	0.008
	(0.020)	(0.020)	(0.018)
Loan value	0.227	0.166	−0.010
	(0.190)	(0.178)	(0.159)
Age	0.256	0.165	0.349
	(0.216)	(0.210)	(0.228)
Income	−0.514	−0.389	−0.015
	(0.500)	(0.476)	(0.442)
GVA change	−0.112	−0.076	−0.070
	(0.168)	(0.164)	(0.169)
Constant	19.552	26.670	5.368
	(17.965)	(19.111)	(16.372)
Controls	Yes	Yes	Yes
Specification	OLS	OLS	OLS
Observations	47,261	47,261	47,261
Adjusted R2	0.002	0.002	0.002
Note:			*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level. Neighbourhood controls include average LTI, house price, change in household networth between House price

7.6 Location Robustness

There might be a few reasons why cash-flow effect of employment differs by region. When I perform my regional heterogeneity calculations I assume that my identified effect is common across the country and only the mortgage and employment structures vary. That seems like a good first pass but it is useful to test how stable my results are when I exclude each of the twelve UK regions from the main regressions.

The first column in Table 18 shows the strongest result I get from dropping one of the regions. The cash-flow shock coefficient is around 40bp higher than the main estimates. At the other end of the spectrum, when I exclude London my identified effect falls by over a quarter. It is not surprising that in a dynamic city like London businesses are very responsive to local demand shocks. Importantly, London is also a very large employer and it has a relatively high share of employment in the *food and drink* sector described above. In fact, within locally non-tradable employment, London employs 23% more people in the *food and drink* sector than the average region.[59]

In the final column I exclude Northern Ireland because its mortgage market looks slightly different to the rest of the UK. In particular, there were many more variable-rate mortgages issued before 2008 so one concern might be that much of the cash-flow shock variation emanated from Northern Ireland. The table shows the main coefficient is moderately lower for Great Britain on its own.

7.7 Cash-Flow Shock Timing

The empirical strategy I use allows for some time to pass between when households receive the cash-flow shock and when employment is expected to respond. That is partly because the employment data are only available every April. The next most sensible specification to run is to assume that spending and employment decisions reacted almost instantly and compare the cash-flow shock to the employment change between April 2009 and April 2008. The first column of Table 20 shows the identified employment effect is very similar to the main estimate for β and strongly statistically significant.

I can also run a placebo test on employment changes in the years before when the cash-flow shock actually occurred. To the extent that there was a high degree of turnover in the mortgage market during the Great Moderation, we should expect insignificant coefficients. The second and third columns of Table 20 confirm the employment-effect estimates are not significant at the 5-percent level.

7.8 Neighbourhood Sample and Growth Rates

When the 400,000 locally non-tradable establishments are assigned their nearest node I am left with just under 50,000 neighbourhoods that span the area of UK. I argue that it makes sense to drop locations with the largest number of employees in the because large establishments (or

[59] The average employment split across locally non-tradable firms is 31% in the *food and drink* sector, 8% in the *vehicle* sector and 61% in the *retail* sector.

Table 18: Regional Robustness

	Annual employment growth, pp		
	(1)	(2)	(3)
	exc. Yorkshire	exc. London	exc. Northern Ireland
Cash-flow shock (% of income)	4.98***	3.01***	3.66***
	(0.82)	(0.85)	(1.14)
Controls	Yes	Yes	Yes
Specification	OLS	OLS	OLS
Observations	43,665	45,963	44,013
Note:			*p<0.1; **p<0.05; ***p<0.01

	Proportion of locally non-tradable employment, %		
	Yorkshire	London	Northern Ireland
Food and Drink	29	38	27
Vehicles	8	4	10
Retail	62	58	63

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level. Lower panel provides employment summary statistics.

groups thereof) have the potential to skew the results of the weighted regressions. In the main specification I therefore drop the largest 0.5% of nodes. The first column of Table 19 shows that when all locations are included in the regression (an extra 300 locations) and the estimate of the cash-flow shock is only slightly smaller and still statistically significant to a high degree of precision. This should provide some reassurance that the cash-flow shock I identify is consistent with the effect experienced at the most dense locally non-tradable employment sites in the country.

The second and third columns of Table 19 show the results for β under alternative growth rate definitions for employment. Column 2 shows that using the conventional growth rate of $\frac{x-y}{y}$ actually leads to almost exactly the same estimate for β. Although the bounded growth rate I use in the main specification is somewhat unusual, it turns out not to make much difference. In the final column I use the log change in employment. Column 3 suggest that a 1% average cash-flow shock for a neighbourhood leads to almost a 10pp increase in the locally non-tradable growth rate and this is statistically significant at the 1% level. It is more than double the estimate of β in the main specification because the log specification has a larger relative error for large negative growth rates (i.e. $\ln(1+x) \xrightarrow[x\to-1]{} -\infty$), which biases the results upwards. Intuitively, the specification gives more credit than it should to the cash-flow shock when an establishment cuts employment back by 80% instead of 90%. Although the log specification therefore provides more concrete evidence of a cash-flow shock and is consistent with other recent establishment-level work (e.g. Giroud and Mueller (2017)), I argue the central specification I use is the most

Table 19: Sample and Growth Definition Choices

	Annual employment growth, pp		
	(1)	(2)	(3)
	Include outliers	Raw percentage	Log change
Cash-flow shock (% of income)	4.44***	4.58***	0.096***
	(1.44)	(1.42)	(0.021)
Controls	Yes	Yes	Yes
Specification	OLS	OLS	OLS
Observations	47,561	47,261	43,731
Note:			*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level.

Table 20: Employment Response Timing

	Annual employment growth, pp		
	(1)	(2)	(3)
	2008-2009	2007-2008	2006-2007
Cash-flow shock (% of income)	5.03***	−4.23*	3.42*
	(1.39)	(2.36)	(1.88)
Controls	Yes	Yes	Yes
Specification	OLS	OLS	OLS
Observations	47,176	47,321	45,007
Note:			*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are clustered at the local-authority level.

helpful for understanding the cash-flow shock at the micro level.

7.9 Standard Error Corrections

7.9.1 Standard Error Clustering by Administrative Region

It is common to cluster standard errrors by contiguous administrative regions, partly because these are often the level at which data are collected. Figure 17 below shows an example. Establishment A serves a neighbourhood defined by a circle with radius r. Establishment B falls outside that neighbourhood but is in the same administrative region. At the same time, establishments C and D are within establishment A's neighbourhood but fall under different regions. The standard approach of clustering at these administrative regions imposes non-zero and zero correlation precisley the wrong way round for these establishments.

In the main regressions I cluster at the local-authority so the results are comparable to the bulk of the outstanding literature. In this Appendix I develop two alternative clustering methods to ensure the significance I find for the main coefficients is not an artifact of the arbitrary local authority boundaries used in the UK.

Figure 17: Stylised Clustering Problem

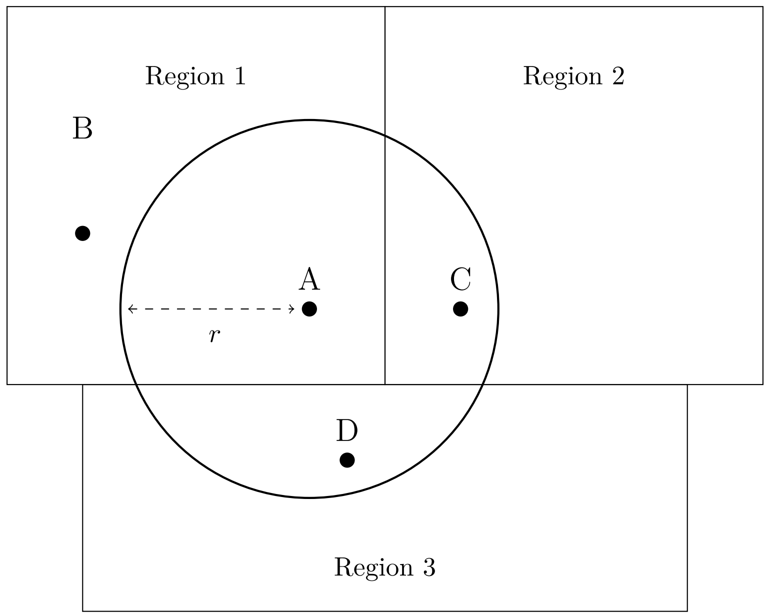

7.9.2 Grid Simulations

Administrative boundaries have two issues that make them inconvenient for using as cluster definitions for my purposes. The first is that they are often the wrong size. In this example, local authorities are heterogeneous in size and therefore do not capture enough of my estimated establishment neighbourhood, or stretch well beyond the reach of the locally non-tradable businesses. The second problem is that administrative areas are arbitrary (or the basis of their formation is orthogonal to the economic behaviour being studied). Both problems can be partly solved by creating a tessellating grid of squares that correspond approximately to the neighbourhood definitions in my baseline regressions. This grid is made up of squares with vertices of length $2r$.

In Figure 18, Box 1 shows the positioning of a grid that happens to line up closest with the approximate neighbourhood of establishment A. Box 1 has establishment A at its centre and the grid can be defined relative to the origin point, O. Establishment A is now allowed to be correlated with establishments C and D but not with establishment B. But choosing a grid structure at the origin O is also arbitrary. I can shift that origin slightly in the lateral and vertical directions to move the entire grid across the country. Box 2 now defines a new cluster where establishments A and C are allowed to be correlated. But there are now other clusters beside and below Box 2 that contain establishments B and D, respectively.

Figure 18: Moving the Tessalating Grid

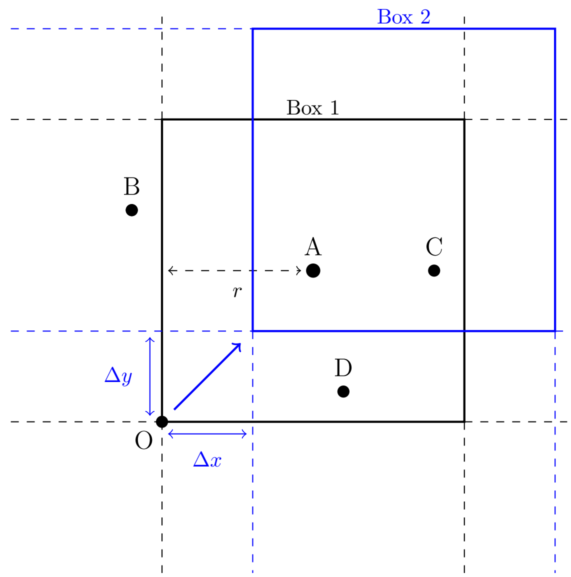

Since all of these grid positions are arbitrary, I am able to perform thousands of simulations to understand how the standard errors vary when clustering in this way. I take the baseline regression as that of cash-flow shock as a proportion of gross income, referred to in Equation 5 in the main text. The histogram in Figure 19 shows 10,000 different simulations for a grid of 20 square kilometres. The domain of the distribution is relatively narrow, which suggests taking the local authority clustering approach is a sensible benchmark. If anything, it is conservative. More than 77% of the simulations produce a standard error below the baseline standard error, and the mean is 3.5% below. This provides reassurance that the significance of the main results is not driven by the unusual partitioning of the microdata.

Figure 19: Clustered Standard Error Simulations

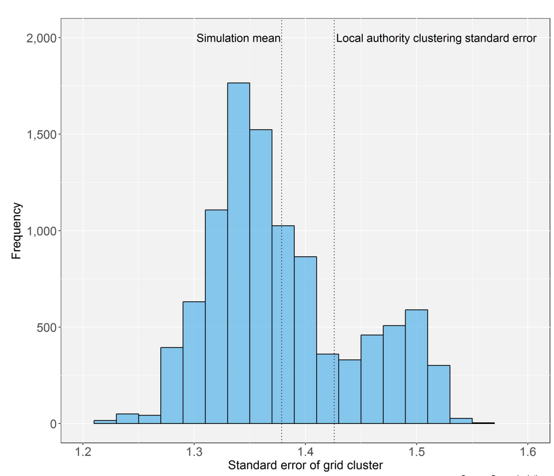

7.9.3 Spatial Clustering

An alternative to using contiguous clustering regions is to spatially cluster the establishments based on the pairwise distance between all observations. We know the standard weighted OLS estimator yields

$$\beta = \left(X'WX\right)^{-1} X'Wy \tag{9}$$

In the special case of identically distributed and independent errors, the variance collapses to the standard expression

$$\hat{V}(\hat{\beta}) = s^2 \left(X'WX\right)^{-1} \tag{10}$$

By relaxing this assumption, the variance can be generalised to

$$\hat{V}(\hat{\beta}) = \left(X'WX\right)^{-1} \left(X'W(\Omega \otimes \mathbf{\Gamma})WX\right) \left(X'WX\right)^{-1} \tag{11}$$

Where Ω is the covariance of the error terms

$$\Omega = \mathbb{E}\left[\varepsilon'\varepsilon|X\right] \tag{12}$$

$$= \begin{bmatrix} \sigma_{1,1} & \sigma_{1,2} & \cdots & \sigma_{1,j} \\ \sigma_{2,1} & \sigma_{2,2} & \cdots & \sigma_{2,j} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{i,1} & \sigma_{i,2} & \cdots & \sigma_{i,j} \end{bmatrix} \tag{13}$$

Then $\mathbf{\Gamma}$ is the matrix of potential correlations between neighbourhoods and is populated by indicator variables

$$\mathbf{\Gamma} = \begin{bmatrix} 1 & \mathbb{1}_{1,2} & \cdots & \mathbb{1}_{1,j} \\ \mathbb{1}_{2,1} & 1 & \cdots & \mathbb{1}_{2,j} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbb{1}_{i,1} & \mathbb{1}_{i,2} & \cdots & 1 \end{bmatrix} \tag{14}$$

At a minimum, we know the diagonal is populated by 1's because heteroskedasticity is permitted. In addition, we can restrict attention to only those neighbourhoods that overlap by enough to be thought of as correlated. At one extreme the cut-off can be zero, which is the same as assuming there is no spatial correlation between establishments. At the other extreme an infinite cut-off is equivalent to clustering at the UK level. We can define different thresholds of overlap to determine whether to include a given $\sigma_{i,j}$, the correlation between two neighbourhoods, in the total variance. This provides a way of implementing spatial correlation in the spirit of Conley standard errors. If the data were aggregated to the local-authority level $\mathbf{\Gamma}$ would be block diagonal and the problem would collapse to simple one-dimensional clustering.

The indicator function takes the following values

$$\mathbb{1}_{i,j} = \begin{cases} 1 & : \text{if } distance(i,j) \leq K \\ 0 & \text{otherwise} \end{cases}$$

The estimated variance is then

$$\hat{V}(\hat{\beta}) = \left(X'WX\right)^{-1} \left[\sum_{i=1}^{N}\sum_{j=1}^{N} w_{i,j} q_i' \hat{\varepsilon}_i \hat{\varepsilon}_j' q_j \mathbb{1}_{i,j}\right] \left(X'WX\right)^{-1} \tag{15}$$

Figure 20 shows the decisions rule for the indicator function. When two firms are at most K apart their employment changes are allowed to be spatially correlated.

Figure 20: Cut-off Distances for Spatial Correlation

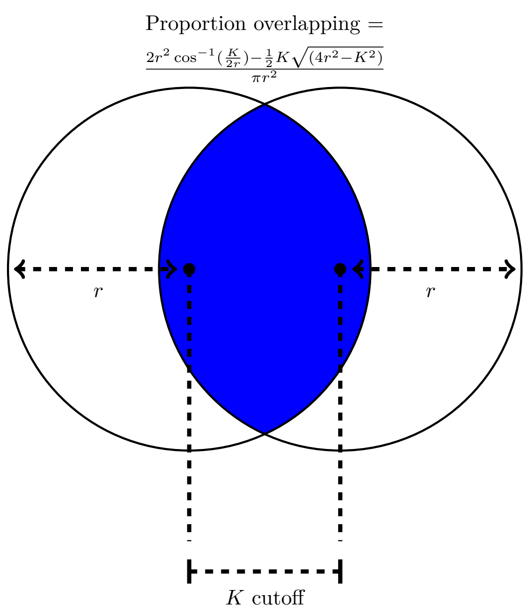

There is a simple mapping between the cut-off distance K and the proportion two neighbourhoods overlap, for $K \in [0, 2r]$. Figure 20 shows the estimated standard errors for different proportion neighbourhood overlaps in the baseline cash-flow shock regression. Two noticeable features spring out. First, one might expect the curve to possess a negative slope because a larger proportion overlap is associated with smaller clusters but this is not what the line shows. Second, for all overlap proportions and cut-offs less than 20km, the standard error is between around 1.3 and 1.5. As before, this robustness check shows that the significance I find in the main regressions is unlikely to be artificially driven by the overlapping nature of my neighbourhoods.

Figure 21: Spatially Correlated Standard Errors