Shocks and Policy over Time and Space

Fergus Cumming

Nuffield College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy in Economics

Trinity 2018

Disclaimer

This thesis 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 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 data sets that may not exactly reproduce National Statistics aggregates.

Abstract

This thesis aims to advance our understanding of the spatial and temporal propagation of economic shocks across the UK economy.

The first chapter sets out the data I will use throughout this thesis. The second examines the role of cash-flow shocks on household and firm behaviour. I find that the reduction in monthly mortgage payments resulting from monetary policy stimulus in the autumn of 2008 significantly boosted employment in the depths of the Great Recession. The neighbourhood-level cash-flow effect varied with the proportion of fixed-rate mortgagors across the country, and its interaction with the structure of local employment markets led to regional heterogeneity in the potency of monetary policy. I find that the effect of monetary policy on overall annual employment growth varied by around 1.5 percentage points across local authorities.

The third chapter explores how local demand shocks associated with house price changes affect employment growth, and whether these shocks are transmitted across firm establishment networks. I find little evidence that credit or informational frictions have played much of a role since the Great Recession. But the employment effects of these local demand shocks do extend beyond administrative boundaries and ripple short distances across the country.

The fourth chapter revisits the Granular Hypothesis and its implications for the UK. I find that the UK firm-size distribution exhibits a quantitatively fat tail, but this is dominated by firms that generate revenues from customers living nearby. The size of such firms is naturally driven by the number of establishments in their network, which means they tend to have low geographical concentration. Since I find limited evidence for establishment-network transmission of local demand shocks, this suggests the propagation of idiosyncratic shocks to large firms is more likely to emanate from supply chain disruptions than from shocks to demand.

Dedicated to my parents.

Tempus fugit.

Acknowledgements

This work would not have been made possible without the support and encouragement of a large number of people. My progress has involved standing upon, stepping between and being propelled forwards by a whole phalanx of giants.

I am especially indebted to my supervisors for their guidance, understanding and tireless work to ensure this work reached its full potential. From the very start, Martin Ellison challenged my thinking and focussed my attention on the important questions. He pushed me, while letting me make my own mistakes, and showed me the value of starting from the right place (and sometimes taking a step back). John Muellbauer's considered words throughout my studies always gave me the confidence to reach for new heights and face the intellectual puzzles with fervour. Together, their perspicacity and vision for my work has been a source of motivation and their kind words of promotion are gratefully acknowledged.

I am also indebted to other members of Nuffield College and the Department of Economics for their generosity of spirit and success in creating an intellectually stimulating environment where ideas can flourish. In particular, I would like to thank Steve Bond, Michael McMahon and Guido Ascari for their advice, in areas relating to their specialities and beyond. I was often halfway down the corridor before I appreciated how profound their ideas had been. My only hope is that they eventually realised the impact their words had on my thinking.

Most of my work would not have been possible without the Bank of England and I am grateful to the institution as a whole for its multi-faceted support during my DPhil. More than that, I am lucky to have had the support of many fantastic colleagues who have given their time in countless meetings to contribute to the success of this endeavour. While that list is too long to detail in full, special thanks goes to Fergal Shortall for his professional support over the past couple of years. I have been positively impacted by all my previous managers but I am especially grateful to Niki Anderson and Lewis Webber who provided me with the opportunity that now counts as my first foray into research almost a decade ago.

In more recent times, I am very much obliged to John Barrdear for his near-weekly beckons to the white board, Angus Foulis for his insights regarding the stories behind the data and Tom Smith for his inspiration regarding a good number of the ideas contained within this thesis. Andy Haldane's demonstrations of the power of ideas were a constant source of drive and working on his speeches proved to be the foundations of much of my research herein. I would go on to list whole divisions across the Bank but their acronyms will surely have changed beyond all recognition for posterity.

I was fortunate to have access to ONS microdata and that was made possible by the enduring efforts of the Secure Research Service in Titchfield. More than anyone else who supported me, Folkert and the team had to put up with my last-minute deadlines, densely-packed requests and glutinous disk-space usage. Without the ONS staff going far beyond the basic service-level agreement, this work would have taken significantly longer. They are all a credit to the public sector's ambition to facilitate research.

At various points along my journey I benefitted from comments and discussions from a number of people who generously gave their time when they almost certainly had more pressing things on their plate. Although the topics of many of these discussions top my catalogue of roads not travelled, I hope that I have many more journeys ahead of me. In Oxford I would like to thank James Best, Andrew Dilnot, Basile Grassi and John Vickers; and in London, Charles Bean, David Baqaee, James Cloyne, Paul Tucker, Nicolas Stern, Ricardo Reiss and David Miles.

I fondly look back at my time in Mission Point and now appreciate the encouragement my friends Olek and Francis gave me to make the journey along the M40. My countless trips back to London (in search of stars more than stardom) were made possible by many accommodating friends. Nick, Tash, Oli, Divya, Adam and Jules deserve a special mention. But those who only look to the past or present are certain to miss the future. I will miss the friends I have made in Oxford, especially Tom, Julian and Jemima, but I know it will be easy to contact them again in years to come, preferably at our regular catch-ups but in the worst case via their public profiles attached to the institutions they will be running. My friends, teammates and breakfast partners at the Triathlon Club kept me sane when my data were misbehaving but they failed miserably in their attempts to convert me to enjoy swimming.

I am very lucky to have had the far-reaching support of my parents and my aunt Dorothy. Trips to Clanfield and Dundee were always a vital dose of academic recuperation, and always a joy. I have been blessed by their advice and optimism for as far back as I can remember but this is a rare opportunity to document it in writing. Thanks also to Anna and Nick who blazed a not-dissimilar trail ahead of me and frequently offered up a new perspective and well-chosen words of advice, as well as their conveniently-located sofa bed.

My final thanks are reserved for Emily, without whom I cannot imagine finishing my DPhil. My decision to go to Oxford was taken before we met, but ever since we have been together, our travels have been shared. Emily is a constant source of love and no words can describe how influential she has been throughout this project. This was a marathon, not just a DPhil, and we did it together.

Contents

List of Tables

List of Figures

Page

Introduction

One of the lessons following the Great Moderation was that small shocks can be amplified to create substantial effects on economic variables that matter to individuals. While banking-sector impairment weighed down on economic output in the Great Recession, it reminded us that the narratives surrounding the propagation of unexpected impulses are often found wanting. This thesis aims to deepen our understanding of the transmission mechanism of economic shocks across time and space.

Sometimes those economic shocks are unwelcome, especially if they are seen to take the economy away from some preferred path. But just because they are (by definition) unexpected does not necessarily mean they are best avoided. In the summer of 2008, fewer than 10% of households expected interest rates to fall substantially in the coming months. Yet fall substantially they did. The decision to reduce policy interest rates by 400bp in the autumn that year offset a number of other adverse shocks and almost certainly shielded the UK economy from some of the more extreme outcomes witnessed in the Great Depression.

After setting out some of the salient features of the data I will use in this thesis in Chapter 1, in Chapter 2 I explore the effect this aggressive monetary easing had on the economy. In particular, I investigate how the reduction in policy interest rates flowed through to employment outcomes at both the neighbourhood and national level. The propagation channel is straightforward: unexpectedly lower monthly mortgage repayments induced a favourable cash-flow shock for mortgagors. Spending increased as a result, which propped up employment. But the interaction of local mortgage and labour markets led to spatial heterogeneity in the effect of monetary policy.

I find that the average household received a windfall equivalent to about 5% of their gross annual income, but the aggregate cash-flow shock was higher in areas where there were more variable-rate mortgages. Some of the additional spending arising from these cash-flow shocks occurred in firms near to where people lived. Regional variation in the distribution of different types of mortgage therefore allows me to identify how employment in these shops, car garages and restaurants responded to neighbourhood cash-flow shocks. I estimate that a 1 percentage point (pp) reduction in interest rates led to annual employment growth for these firms increasing by around 3.5pp.

By using data accurate to a few hundred metres, I am able for the first time to go beyond quantifying the impact of monetary policy for the economy as a whole. Not only does the proportion of variable-rate mortgages vary across space, but so does the make-up of local labour markets. I show that the employment effect of monetary policy was greatest in parts of the country where there was strong pass-through from interest rates to mortgage payments, and where the local-facing economy made up a large fraction of employment. My results suggest that the impact from the change in policy interest rates on one-year total employment growth in the Great Recession varied by around 1.5pp across local administrative regions. To the extent that mortgage and labour market characteristics evolve over time, so too will the spatial traction of monetary policy.

In Chapter 3 I move away from explicit policy and explore the employment effect of local demand shocks proxied by house price changes over the past decade. Changes in collateral values had less impact on debt servicing costs but influenced the amount households could borrow, as well as affecting their overall wealth. I show that increases in house prices were associated with meaningful increases in employment growth at the same firms that relied on customers living close by in Chapter 2.

These shops, car garages and restaurants often belong to chains and I can use these business networks to test how local demand shocks spread across the country. One hypothesis is that shocks to part of a firm's network might affect employment decisions in other parts of the chain if the firm is reliant on access to credit markets. Another hypothesis is that individual shops might use information about the demand prospects gleaned from sister establishments to inform how they should prepare for the future. The final hypothesis is that these networks of firms are of limited importance for the spread of local demand shocks.

I find that in the years since the Great Recession there is little evidence that business networks have played a significant role in the propagation of local demand shocks. The most convincing reason that demand shocks do ripple across the country is that people travel across administrative borders. That naturally limits the scope for economy-wide spatial propagation.

In Chapter 4 I take a step back and look at UK firms as a whole. Like most economies, employment across UK firms is far from evenly distributed. This has always been true. Powerful institutions such as the East India Company in the 18th Century, British railway companies in the 19th Century or large supermarkets today have often employed substantial fractions of the labour force. At the other end of the spectrum, there have always been a large number of very small firms that, together, make up a relatively modest share of total employment.

The presence of very large firms could have implications for the propagation of shocks in the economy. For example, an idiosyncratic shock that hits a large employer could impact the rest of the economy if both the shock and firm were large enough. This Granular Hypothesis holds intuitive appeal and has echoes of the *too-big-to-fail* problem in the banking sector following the Great Recession. But away from financial crises, it is less obvious why large firms in the real economy should be a source of vulnerability for the economy as a whole. Should we be more worried about shocks to demand or shocks to supply for the largest firms?

I show that some of the largest firms in the UK are the shops, car garages and restaurants I have focussed attention on in the previous chapters. These firms tend to be large because they have geographically expansive networks. But in Chapter 3 I show there is limited evidence for network propagation of demand shocks, which suggests a key source of vulnerability for the largest employers is likely to be supply chain disruptions.

Chapter 1

Geography, Firms and House Prices

This chapter provides a brief outline of the data sources used in the later analysis. Much of the literature studying the United States aggregates data by state, county or zip code. In the UK the most sensible unit of aggregation is the local authority, for which a great deal of data are collected in the census and through other statistical surveys. The Office of National Statistics (ONS) started publishing monthly house price index data for the complete list of local authorities at the start of 2005, though more sparse data exist back to 1995. Establishment-level employment data come from the confidential UK business register, which includes all firms subject to value-added tax (VAT).

1.1 Geographic Data

The United Kingdom is split up into 389 local authority districts (henceforth local authorities), which correspond to the fourth-level subdivision of countries following the European Union's Nomenclature of Units for Territorial Statistics (NUTS).[1] Each local authority is referred to in a particular way depending on where it is located in the UK. England contains a mixture of 325 *Unitary Authorities* and *Local Authority Districts* (including *Metropolitan Districts* and *London Boroughs*); Wales has 22 *Unitary Authorities*; Scotland has 32 *Council Areas*; and Northern Ireland has 11 *Local Councils*. Institutions such as the ONS provide good data coverage at the local-authority level via the census and other surveys so they are often the natural geographic unit to study local demand shocks. Local authorities are grouped into twelve regions, for which a great deal of data are available for understanding broad differences across the UK.

Local authorities vary greatly in their economic and geographic characteristics. The Highland region of Scotland covers 26,000 square kilometres compared to the 3 square kilometres of the City of London. At the same time, Birmingham has a residential population of 1.1 million people compared to the City of London's 10,000. This variation across administrative districts poses a challenge to the identification of local demand shocks but it is not uncommon in studies of these nature. In the US, the population of the largest county is 100,000 times that of the smallest.[2] The average population of a UK local authority is around 130,000

[1]For the purposes of this study I have dropped the Isles of Scilly, which brings the total to 390.

[2]Los Angeles Country has around 10 million inhabitants compared to Loving County's 95 in

people, which is around 20% higher than the average US county.

The UK is further subdivided into around 10,000 administrative *wards* but relatively little data exist for these beyond social trends captured by the decennial census. There are a similar number of Middle-layer Super Output Areas (MSOAs), which are well suited to granular regional studies in principle as they are formed of relatively homogeneous postcodes, though MSOA-level data are collected relatively infrequently and are often of limited scope. MSOAs were created after the 2001 census and each has a mean population of a little over 7,000 people.

The geographic divisions described above are a natural starting point when exploring regional variation (e.g., Chapter 3). Grouping similar areas together helps parse out various factors that determine economic behaviour. But the ideal identification methodology recognises that no set of pre-defined administrative or statistical boundaries fully reflects the economic transactions of interest. For that reason, I develop a spatially accurate identification strategy in Chapter 2, which can be thought of as an extension to the MSOA classification used by the ONS.

1.2 Firm Data

The main data set I use in the following chapters 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 (henceforth establishment) level. Since

Texas.

the IDBR is based on Her Majesty's Revenue and Customs (HMRC) tax data, it captures close to the universe of economically active firms in the UK. 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. Enterprises are made up of one or many establishments, which are physical workplace locations 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 and business demographic information for both. Throughout this thesis, I use the terms enterprise, firm and business interchangeably.

To identify the effects of cash-flow and other local demand shocks, this paper follows Giroud and Mueller (2017b) in categorising firms into *locally non-tradable*, *locally tradable* and *other* firms.[3] The original classification of Mian and Sufi (2014) was 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. My identification of local demand shocks for the UK economy requires a mapping from the four-digit North American Industrial Classification System (NAICS) codes used in the US to the so-called two-digit Standard Industrial Classification (SIC) system used in most of Europe. While the exact mapping between the two systems at a very granular level at different points in time is far from straightforward, the relatively

[3] Mian and Sufi (2014) further subdivide *other* into construction and non-construction firms. I adopt the term *locally* non-tradable and tradable to emphasise the spatial and within-country focus of my analysis.

high-level categories used for this analysis in the time period of interest match up almost exactly.[4] The 25 locally non-tradable and 82 locally tradable four-digit NAICS-12 industries therefore map into 13 and 66 SIC-03 groups, respectively, for this study.

Appendix A shows the exact mapping between the NAICS-12 and the SIC-03 systems I use in this thesis. Previous work in the UK has explored the geographical concentration of various industries (e.g., Campos (2012)). Of the 30 least geographically concentrated industries in the UK, over half relate to non-public and non-construction activities. The locally non-tradable definition used in the rest of this analysis captures the majority of the remaining industries. Wholesale activities and transport systems are excluded because they are unlikely to capture local demand effects convincingly.

Of the 30 most geographically concentrated industries in the UK, about half are captured in the categorisation of Mian and Sufi (2014); 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 locally tradable in that the firms do not garner most of their sales from nearby customers.

According to the primary classification system used in this study, locally non-tradable employment made up just over a fifth of UK aggregate private-sector employment in 2017, which is very similar to that in the US. Of the 13 locally non-tradable industry groups, the largest employers are retailers of groceries, restau-

[4] For example, the NAICS-07 category of *Automobile dealers* (4411) maps into the SIC-03 group of *Sales of motor vehicles* (501).

rants and other general merchandise stores. Data from the National Transport Survey (see Department for Transport (2016)) suggest that local 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.[5]

1.3 House Price Data

The ONS began publishing monthly house price indices for local authorities in the first half of 2016. Price indices covering local authorities in the whole of the UK are available from January 2005, though exhaustive data for England and Wales are available back to 1995. The time series in Figure 1.1 shows the remarkable evolution of house prices in recent times. Between 1995 and 2005 house prices in England increased threefold. In the following ten years, English house prices increased by a mere 30%. These aggregate trends, however, hide both the temporal and spatial variation in house price changes that I will exploit for the analysis in Chapter 3.[6]

The few years surrounding the Global Financial Crisis of 2008 (henceforth, Crisis) were especially important for the evolution of house prices across regions.

[5]This compares with the average commute within England for work purposes of a little more than 32km. Distances have been increasing marginally over time 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.

[6]The aggregate profile in Figure 1.1 is also consistent with the Halifax and Nationwide published series, the two other main sources of house price information in the UK. Those institutions use privileged mortgage data to compile their indices but do not produce local-authority indices.

Figure 1.1: National House Price Evolution

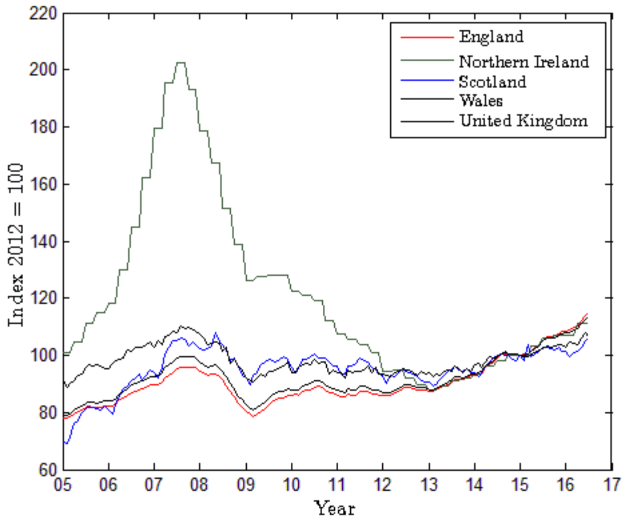

Source: ONS.

Figure 1.2 shows the percentage change in house prices for the UK between 2006 and 2007, and again for 2008 to 2009. The figures show that Northern Ireland experienced the largest short-term gains. In particular, Armagh, Banbridge and Craigavon experienced house price increases of over 60% between 2006 and 2007. As Northern Irish house prices plateaued the following year, local authorities just outside Greater London such as Sevenoaks, Spelthorpe and Windsor continued to experience mid-double-digit percentage annual price increases.

Between 2008 and 2009, house prices fell by around 15% across the UK as a whole and more than double that in Northern Ireland. Nationwide, the correlation between house price changes in contiguous local authorities was around 75% but that figure falls very rapidly as distance between local authorities increases. In Chapter 3 I make use of the fact that regional house price cycles were slightly out of sync with each other since the Great Recession.

Figure 1.2: Spatial Heterogeneity in House Price Changes

(a) 2006-2007

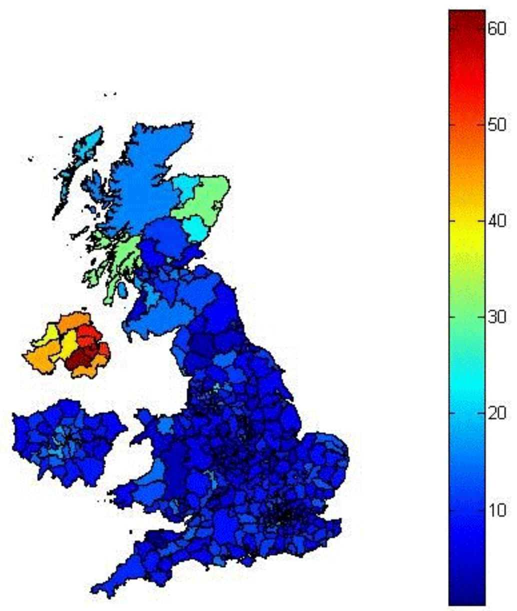

(b) 2008-2009

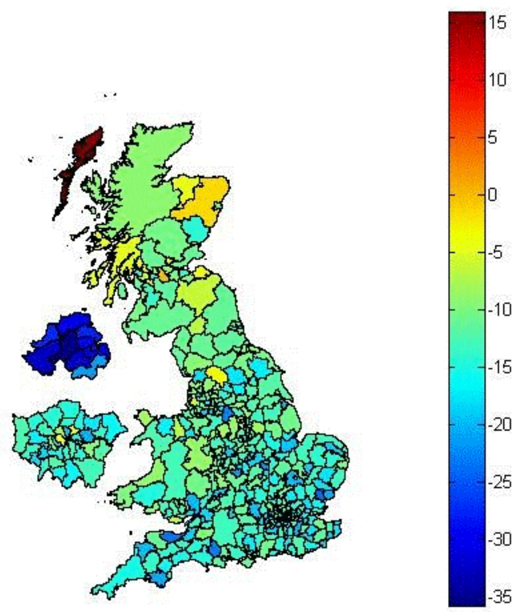

Source: ONS. Figures show the percentage difference between local-authority house-price indices from January to January. Greater London is shown inset below Northern Ireland.

Chapter 2

Monetary Policy and its Effect on Regional Employment

This chapter quantifies the regional impact of monetary policy by combining novel micro datasets with near-universal coverage of UK mortgages and employment. I estimate that on average in the Great Recession a one percentage point fall in interest rates led to a 3.5 percentage point increase in employment growth in locally non-tradable businesses the following year. But the spatial distribution of mortgage and labour market structures led to significant heterogeneity of this effect across the country. I show that the overall effect of accommodative monetary policy on total employment growth in 2010 varied by around 1.5 percentage points across regions.

2.1 Introduction

It is well recognised that monetary policy has redistributionary consequences 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 of monetary policy 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 lead to regional heterogeneity. 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 effect of monetary policy 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 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 simultaneously slackened liquidity constraints and increased permanent income to the extent the shocks were viewed as persistent. In the spirit of Mian and Sufi (2014), I define locally non-tradable firms as those generating revenues from nearby customers and use employment changes to proxy for otherwise unavailable granular consumption data. By modelling the evolution of almost six million mortgage payment flows, I use spatial variation in cash-flow shocks to investigate the response of employment to monetary policy. But 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.[1] I go on to estimate the employment effect of monetary policy at various levels of aggregation.

In the absence of randomised helicopter drops, natural experiments can help us trace out the transmission of cash-flow shocks to aggregate variables. Two features of the UK mortgage market aid the quest for this identification. First, many people were as-good-as randomly assigned a fixed or variable-rate mortgage before October 2008. The average length of a pre-Crisis UK 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

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

on an anticipation of future 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 mortgagors took out both fixed and variable-rate contracts in the immediate run-up to the 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 (see Figure 2.5). This stands in stark contrast to the US, where remortgaging spiked up following the monetary easing.

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. At the same time, around a quarter of mortgagors 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 remains almost unchanged when the

[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.

interest rate wedge between contract types at origination is used as an instrument for the proportion of variable-rate mortgagors in the neighbourhood. But 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.

The distribution of mortgagors on different types of contract was a key driver for the spatial distribution of cash-flow shocks. By aggregating cash-flow shocks to the neighbourhood level I show that, using cross-sectional variation, a 1 percentage point (pp) accommodative monetary policy shock was associated with a mortgagor cash-flow shock of around 75 basis points (bp) and around a 3.5pp increase in the annual employment growth for locally non-tradable establishments. 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. This chapter therefore provides a micro empirical estimate for the aggregate effect of monetary policy on employment and is consistent with, though slightly higher than, standard macro estimates for monetary policy shocks (e.g., Romer and Romer (2004) and Cloyne and Hürtgen (2016)).[3]

The main result of this chapter is to show that the joint spatial distribution of

[3] Applying cross-sectional identification to an aggregate effect requires some strong assumptions, many of which are laid out in Nakamura and Steinsson (2017).

mortgage and labour market structures led to significant heterogeneity in the traction of monetary policy across the country in the Great Recession. Regions where there were a large number of variable-rate mortgagors were particularly sensitive to the direct effects of monetary policy. This was compounded in areas where the local economy was relatively closed-off 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 monetary policy on regional 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. 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.

This chapter 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 me to identify 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).[4] In particular, my results suggest an endogenous link between Auclert's *earnings heterogeneity channel* and the *interest rate exposure channel*. In a similar spirit to Beraja, Fuster, Hurst, and Vavra (2017), I am able to track which regions are most responsive to interest rates in the short run and go on to capture the knock-on effect monetary policy has on locally non-tradable employment. Through this lens, monetary policy drives, as well as responds to, regional employment inequality via changes in local consumption.

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. 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 utilise some

[4] See Calza, Monacelli, and Stracca (2013) and Corsetti, Duarte, and Mann (2018) for studies examining heterogeneous effects of monetary policy through the housing channel.

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 interest-rate-wedge 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 to employment, so is the make-up of local labour markets.

My work also 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)).[5] 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 and

[5] 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)).

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.

The rest of this chapter is organised as follows. Section 2 describes the mortgage data I use in addition to the data set out in Chapter 1. Section 3 sets out the empirical strategy and clarifies the identification approach that runs throughout the whole chapter. Section 4 presents the main results and Section 5 contains some important robustness checks. Section 6 sets the results in context and concludes.

2.2 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).[6] It contains a wealth of information on property, borrower and lender characteristics at the time of origination. The first step is to take these mortgage flows and 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 B goes through the steps needed to transform the mortgage flows into the stock

[6]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.

used in this thesis.[7]

The 5.8m mortgages in my estimated stock form a large and representative sample of the residential mortgage market in 2008. There are no consistent estimates of the true universe of mortgages at this time but 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.[8] Regulatory data for 2015 shows that there were around 6.7m first-charge residential mortgages and Figure B.1 of Appendix B provides an illustration of which mortgages I estimate may be missing from my stock.

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).[9] Mortgage terms tended to be between 25 and 30 years but the periods governing the path of the interest rate (henceforth, contractual maturity) have 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.

[7]From 2012 the FCA began publishing the true stock of mortgages as part of the PSD making the process outlined in Appendix B much more straightforward for future studies. Data quality are increasing over time.

[8]See https://www.cml.org.uk/industry-data/key-uk-mortgage-facts/ for latest estimates.

[9]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.

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 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 my sample were in some way different to the others, biasing my results. In fact, because this study uses cross-sectional variation, bias is only likely to arise if the mortgages missing from my sample 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 2.1.

Figure 2.1: Distribution of Mortgages in July 2008

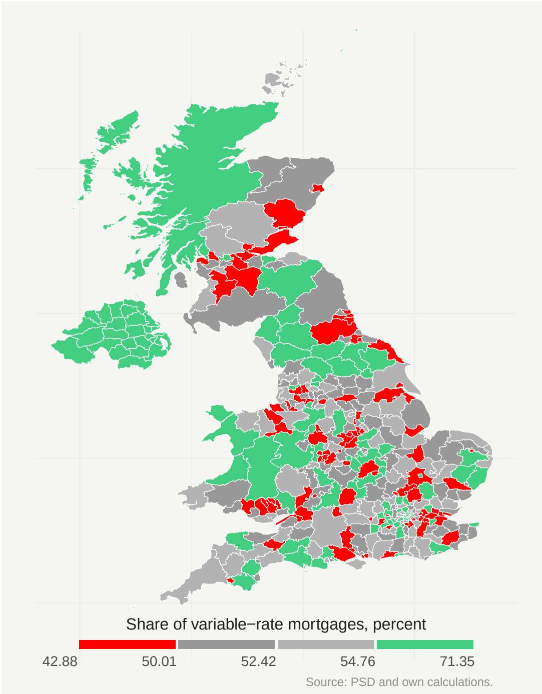

Colour breaks denote quartiles. 389 local authorities shown. Variable-rate mortgages include those past their contractual maturity and linked to the SVR.

I next take my estimated stock and model how the associated characteristics evolved since 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 varies 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 for the cross-sectional variation I use than the level. Appendix B 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 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 2.2 shows the evolution of Bank rate, the two most common (new) mortgage contracts and the SVR.[10]

[10]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 2.2 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.

Figure 2.2: Mortgage Interest Rates

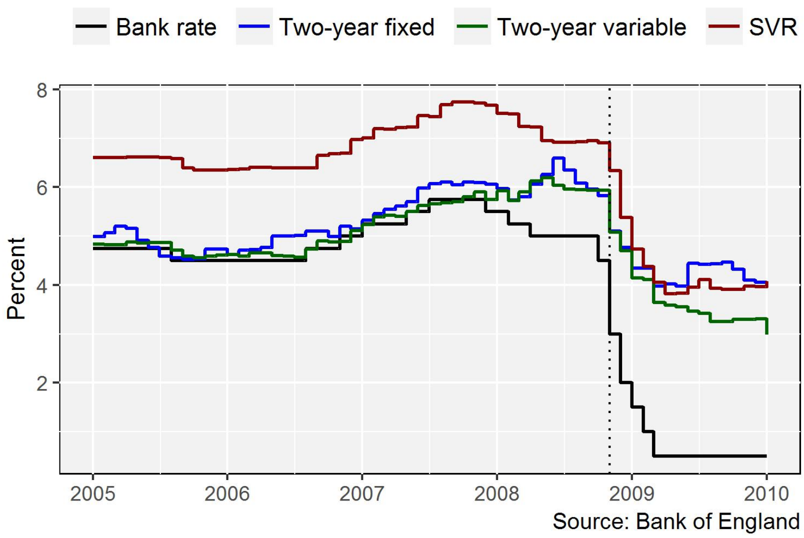

The first series shown is Bank rate. From the Bank of England quoted rates series (with identifying code in brackets): 2-year fixed 75% LTV mortgage (IUMBV34), 2-year variable 75% LTV mortgage (IUMBV48) and the Standard Variable Rate (IUMTLMV).

2.3 Research Design

2.3.1 Central Specification

My empirical strategy tests to what extent monetary policy supports consumption and its knock-on effect to local employment.[11] 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, both in the UK and across the globe. But even if all households benefitted from the support to asset values and their net wealth, only those households on variable-

[11]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).

rate mortgages benefitted from the immediate decrease in mortgage payments.[12] This cash-flow effect was dramatic and we would 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.[13]

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 2.3, 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.[14] 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. Finally, I compare the proportion of variable-rate mortgages and the neighbourhood cash-flow shocks

[12] 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 but the difference was less sharp.

[13] 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.

[14] 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.

to the change in establishment-level locally non-tradable employment in the next financial year, between April 2009 and April 2010.[15]

Figure 2.3: Empirical Strategy Timeline

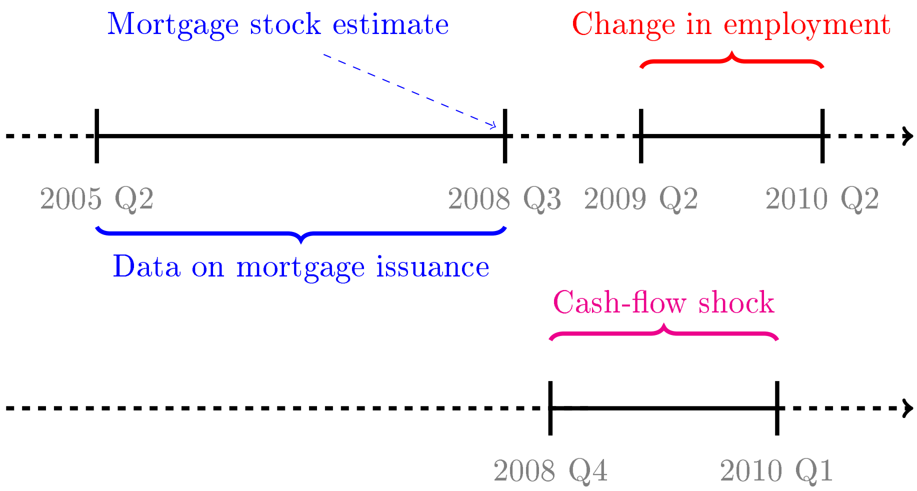

The figure shows the timeline for the central specification. I construct an estimate of the stock in the early Autumn 2008 then calculate the income shock associated with lower mortgage payments for every mortgagor over the following five quarters. The dependent variable is the percentage change in employment 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.

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.[16] Many neigh-

[15]See Appendix C for robustness checks regarding the timing of the employment response to the cash-flow shock.

[16]This means that the centre of the closest neighbourhood is at most around 350m away from

bourhoods 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 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 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.[17]

the true location of an establishment.

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

$$e_i^{10-09} = 2 \times \frac{E_{i,2010}^{NT} - E_{i,2009}^{NT}}{E_{i,2010}^{NT} + E_{i,2009}^{NT}} \tag{2.1}$$

My initial specification is then

$$e_i^{10-09} = \alpha + \beta \times V_{j_i} + \gamma \times X_{j_i} + \varepsilon_i \tag{2.2}$$

The main explanatory variable is V_{j_i}, which captures the proportion of mortgaged households opting for variable-rate contracts in neighbourhood j_i surrounding the establishment. 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 referred to in Chapter 1 as a baseline. In the main specification, j_i 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.[18] A sample of nodes and neighbourhoods is shown in Figure 2.4. 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.

[18] The average US county is a little over 2,500km^2.

Figure 2.4: Oxford Sample Neighbourhoods

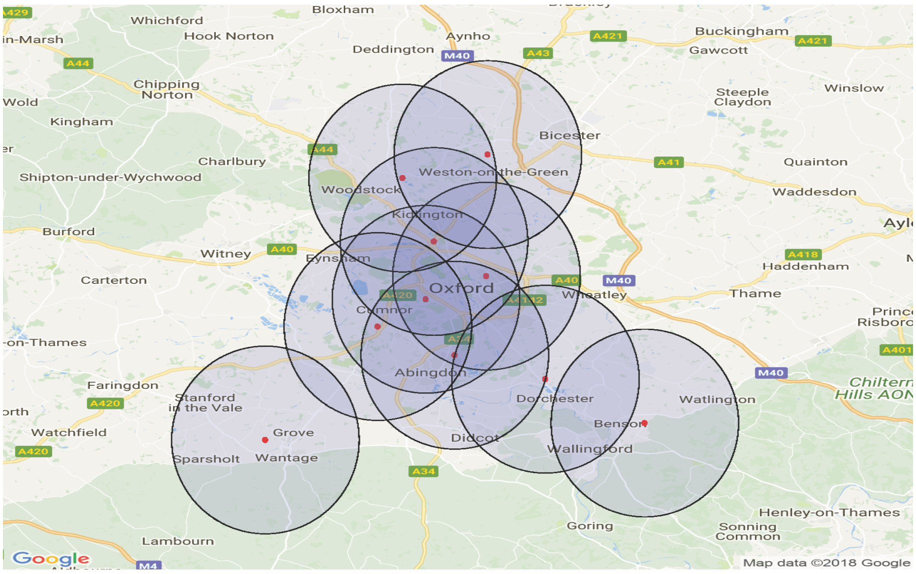

Source: ONS and own calculations. This figure shows a selection of locally non-tradable nodes and neighbourhoods near Oxford assuming the baseline 10km radius.

X_{j_i} 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.[19] These controls use the estimated stock of mortgages, so accurately capture the characteristics of mortgagors in each ring. For more information, see Table 2.3. 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.[20] In

[19] 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.

[20] 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

Equation 2.3 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_i is the sum of these between the start of 2008Q4 and the end of 2009Q4, averaged across all mortgagors in the neighbourhood.[21]

$$\overline{Y}^{\pounds}_{j_i} = \frac{1}{M_{j_i}} \sum_{m}^{M_{j_i}} \sum_{t=2008Q4}^{T} p_{m,t} - \tilde{p}_{m,t} \quad \forall m \in j_i \tag{2.3}$$

In Equation 2.4 the cash-flow shock is scaled by mortgagor annual gross income as of July 2008.[22]

$$\overline{Y}^{\%}_{j_i} = \frac{1}{M_{j_i}} \sum_{m}^{M_{j_i}} \sum_{t=2008Q4}^{T} \frac{p_{m,t} - \tilde{p}_{m,t}}{Y_{m,t_1}} \quad \forall m \in j_i \tag{2.4}$$

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

$$e_i^{10-09} = \alpha^{\pi} + \delta^{\pi} \times \overline{Y}^{\pi}_{j_i} + \theta \times X_{j_i} + \varepsilon_i \quad \forall \pi \in \{\pounds, \%\} \tag{2.5}$$

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

2008Q4 and the end of 2009.

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

[22] 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.

growth.[23] A £1,000 cash-flow shock is associated with a $\delta^{£}$pp increase in locally non-tradable employment growth. All regressions are weighted by establishment size.

2.3.2 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.

2.3.2.1 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 should lead to downward bias in my estimate of β and δ because richer households are more likely to have a lower

[23] 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).

marginal propensity to consume.[24] In fact, 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 are remarkably similar.

Mortgages are poorly understood by households and each mortgage contract has a vast number of non-price terms and options.[25] 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 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 mort-

[24] 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).

[25] 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.

gage 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.[26] I use the population of remortgagors in my sample to explore this. The top panel of Table 2.1 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 quasi-random assignment is also consistent with the observable characteristics of the mortgagors reported in Table 2.2, 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.[27] The latter, espe-

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

[27] Bacon and Moffatt (2012) find that those on fixed-rate contracts were likely to borrow more

cially, might be cause for concern if it means that those on fixed rates are likely to be more sensitive to counterfactual cash-flow shocks.[28]

Although the median values in Table 2.2 suggest that there might be some inherent difference between the types of people opting for different mortgages, there a three grounds for reassurance.

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 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 2.2). Focusing on the difference between the median or the mean is likely to overstate the behavioural heterogeneity of the two groups.[29]

between 1992 and 2001, which is not the case in the PSD data.

[28]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.

[29]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.

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.[30]

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.[31] 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. In the robustness checks I use the wedge between fixed and variable rates as an instrument and it supports the validity of the above identification; my estimate of β remains almost unchanged after instrumenting.

2.3.2.2 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 mort-

[30] 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, 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.

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

Table 2.1: Evidence for Quasi-Random Mortgage Selection

Expected | no relationship, percent

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

Note: Pr*(Fixed) = 63%*

Observed, percent

First mortgage	**Second 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, 66% 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.

gage 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 2.5 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.[32] This subdued refinancing

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

Table 2.2: 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.

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 contract took advantage of the lower rates.

Figure 2.5: Mortgage Approvals

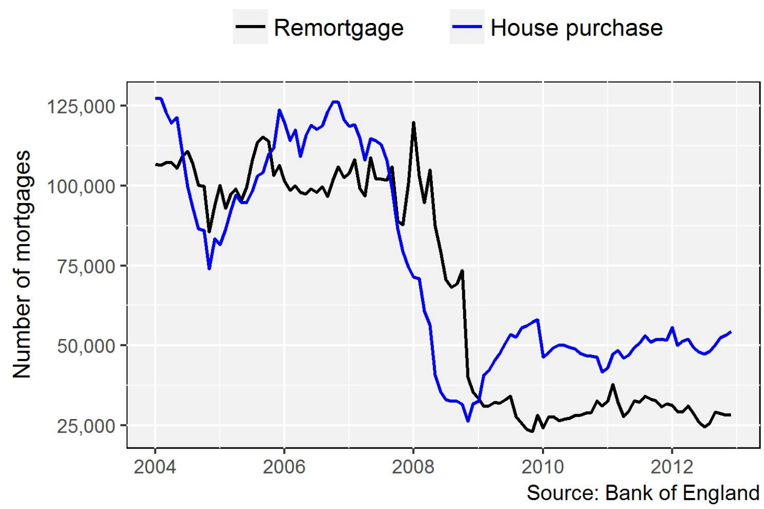

The series correspond to the seasonally adjusted data for housing remortgage (B4B3) and housing purchase (VTVX) from the Bank of England Statistical Interactive Database.

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.

2.4 Results

2.4.1 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 2.6 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 of those 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 2.6: Individual Cash-Flow Shock Distributions

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

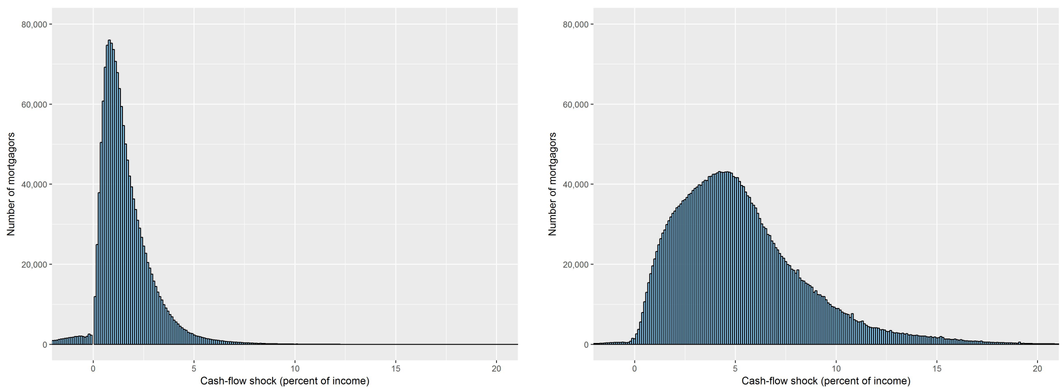

Source: PSD. A total of 4.5 million mortgages are captured in these distributions. 1.2 million mortgages very close to 0 have been dropped from the fixed-rate distribution to make the charts easier to compare. The small number of negative cash-flow shocks occur when people revert or remortgage onto a 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).[33] 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 2.7.

[33] See ONS data series NRJS. This is also consistent with survey evidence. For example, 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 2.7: Neighbourhood Cash-Flow Shock Distributions

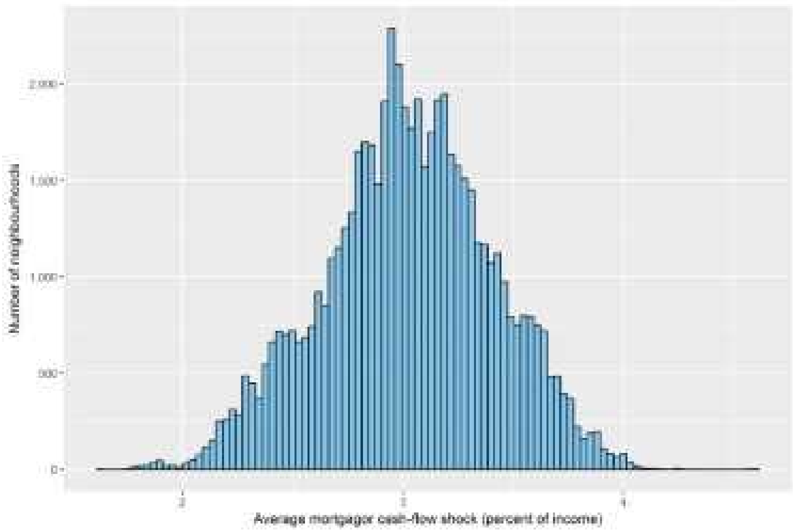

Source: PSD. The average cash-flow shock for mortgagors in a neighbourhood is a function of, among other things, the proportion of people entering fixed and variable-rate contracts, average household leverage and the exact timing of the contracts taken out.

Table 2.3 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.

Table 2.3: 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, 2009	47,261	60.7	222	3	26
Change in LNT employment, 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.

2.4.2 Aggregate Cash-Flow Effect

In this chapter 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 2.4 shows 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 gives us an insight into how monetary policy pass-through is affected by the structure of the mortgage market.

The second two columns of Table 2.4 present my estimates for δ in Equation

Table 2.4: Results Summary

	Annual employment growth		
	(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

2.5 using definitions in Equations 2.3 and 2.4. 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.[34] 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 2.5.

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.[35] 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

[34]The choice of clustering standard errors at the local authority is conservative. One alternative is to construct hundreds of slightly different square grids to makes sure establishments on the boundary between administrative regions are not driving the results. I find that although there is naturally some variation in standard errors, when using these tessellating grids, the average standard error for each variable is almost identical to the main set of results. Another alternative is to use Conley standard errors, which allow for clustering by distance. These standard errors are lower in all cases and so I use local-authority clustering for the remaining analysis to err on the side of caution.

[35]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 2.3.

results therefore suggest locally non-tradable employment could have shrunk by more than 15 percent absent policy intervention.

2.4.3 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. In the UK the most natural administrative region to illustrate my results is the local authority. Taking the full 400bp of policy easing, Figure 2.8 shows the estimated impact of the accommodative monetary policy shock on all 389 local authorities' total employment growth by quartile.[36] Areas shaded green (red) correspond to parts of the country where overall employment was most (least) affected by monetary policy. But the spatial distribution of the total 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 in Figure 2.1. 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.

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

Table 2.5: Main Specification

	Annual employment growth		
	(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
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.

Figure 2.8: Estimated Impact of 400bp Cut in Policy Interest Rate

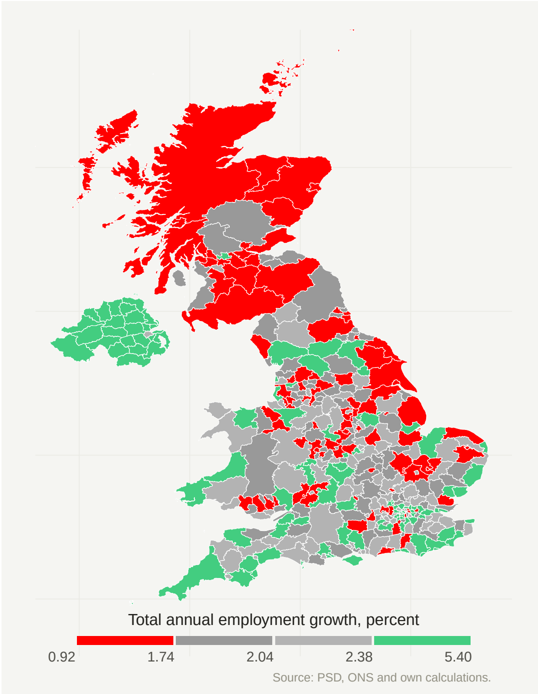

This figure is constructed by taking the establishment-level cash-flow semi-elasticity of employment and combining it with cash-flow shocks and locally non-tradable employment shares for each local authority. Colour breaks denote quartiles.

Table 2.6 provides more details on the distribution of the estimated overall effect of monetary policy across regions. Column 1 combines $\hat{\gamma}^{\%}$ 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 1.55 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, $\hat{\gamma}^{\pounds}$, 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 averaged-size 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.

2.5 Robustness

2.5.1 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 2.9 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 diffi-

Table 2.6: Estimated Heterogeneous Effect of Monetary Policy

Local Authority Percentile	Annual total employment growth, pp: %-shock Observed	%-shock Counterfactual	£-shock Observed	£-shock 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.

cult 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. 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.[37] Even if households

[37]Evidence presented in Agarwal and Mazumder (2013) suggests low mathematical ability is

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.[38] 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.

associated with poor financial decision making. My identification relies more on the level of overall mortgage market comprehension.

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

Figure 2.9: Mortgage Issuance

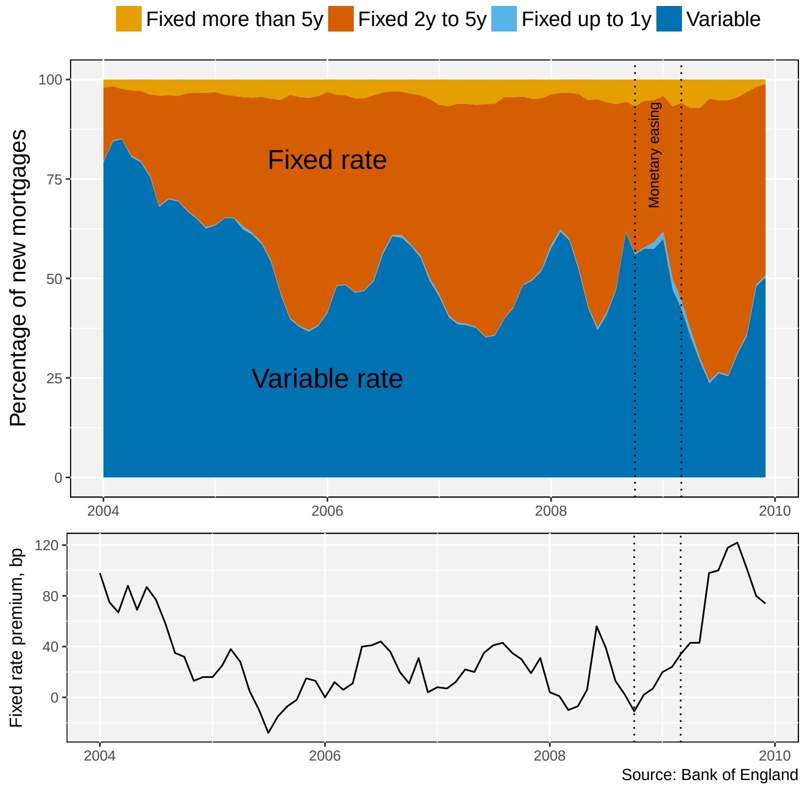

The top panel of this figure shows the variation in mortgage-type issuance across time. The bottom panel shows the time variation in the fixed-rate premium. These data are taken from Bank of England aggregated data and therefore have a slightly different make-up from the rest of the analysis (e.g., they 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 this rate, which is driven by changes in the steepness of the short-end yield curve. To construct my instrumental variables estimate I follow a three-stage least squares approach.[39] The first stage requires constructing a linear probability model to predict the probability that an individual will take out

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

a variable-rate mortgage, P_m, based on the average FRP prevailing in the three months before the mortgage was issued and a host of other individual controls, X_m, including LTI, house price, age and mortgage term.[40] These were selected by testing various specifications in an attempt to achieve a good yet parsimonious fit.

$$P_m = \alpha + \pi \times FRP_t + \lambda \times X_m + \varepsilon_k \tag{2.6}$$

The results from this first stage are shown in the first column of Table 2.7. 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_i. It yields the expected number of fixed-rate mortgages at origination for each neighbourhood.

[40] 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.

$$\hat{\mathcal{P}}_{j_i} = \frac{1}{M_{j_i}} \sum_{m=1}^{M_{j_i}} \hat{P}_m \quad \forall m \in j_i \tag{2.7}$$

I am now in a position to use the aggregated fitted values in Equation 2.7 as an instrument for the proportion of variable-rate mortgages in my neighbourhoods. Figure 2.10 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 2.10: Instrument Relevance

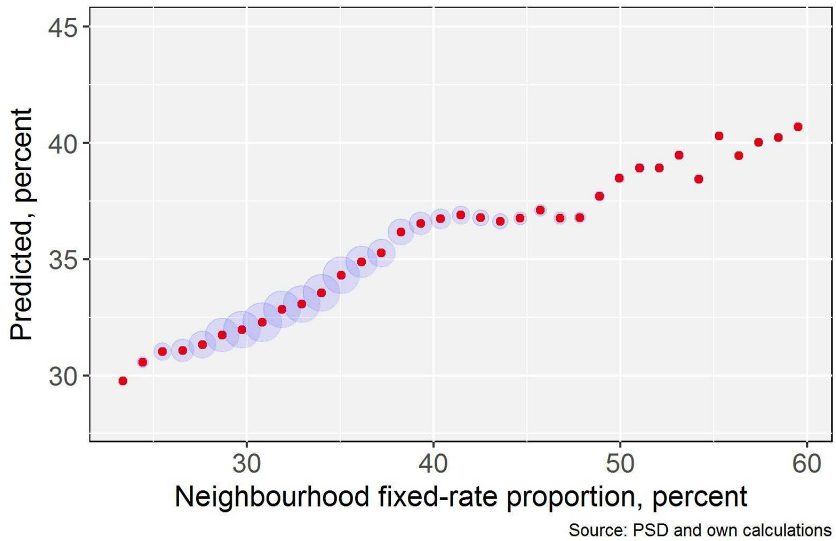

This figure plots the neighbourhood proportion of fixed-rate mortgagors in the summer of 2008 against the predicted proportion of mortgagors taking out fixed-rate contracts after aggregating individual fitted probabilities (in Equation 2.7). Blue discs represent the number of mortgages in each bucket.

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 I am using the fitted values 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 2.7 shows the instrumental variables estimate, which is very close to the main coefficient presented in Tables 2.4 and 2.5. This provides some reassurance that the main specification has been well identified. In particular, it suggests that although there might be selection bias in the distribution of fixed and variable-rate mortgagors, the observable characteristics of mortgagors do a good job of controlling for it and the estimate of β can be interpreted as causal.

2.5.2 Types of Firm

The results in Table 2.5 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 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 my placebo test on establishments that fall under the locally tradable industrial classification outlined in Appendix A. Table 2.8 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

Table 2.7: Instrumental Variables

	Mortgagor level P(fixed-rate)	**Neighbourhood level** Annual employment growth
Fixed-rate premium (bp)	-0.139^{***} (0.001)	
Variable rate (% of mortgagors)		0.178^{***} (0.055)
LTI	2.912^{***} (0.019)	-12.854^{***} (2.939))
House price	-0.024^{***} (0.000)	-0.032^{**} (0.013)
Age	-0.136^{***} (0.003)	0.250^{**} (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^{***} (0.169)	19.660^{***} (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.

worth probing. Firms are asked to report the industrial code that best reflects their activities but this is often prone to error or just difficult to do for a firm engaging in multiple business activities. In the first three columns of Table 2.9 I split my locally non-tradable establishments into three sectors.[41] 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 (see Chapter 4), 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); Kleiner (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 2.9. Chain stores are defined as belonging to enterprises that have more than one es-

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

tablishment. Large establishments are defined as those that employed more than 25 people in 2009.[42] 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.[43]

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

[43]The definition of small establishments in the Table 2.9 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 2.8: Locally Tradable firms

	Annual employment growth		
	(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 2.9: Firm Types

	Annual employment growth					
	(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.

2.5.3 Neighbourhood Definitions

I argue that one of the contributions of this chapter 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 2.10 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 (see Chapter 1). 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.[44] 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 chapter.

[44]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 2.10: Neighbourhood Specifications

	Annual employment growth					
	(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.

2.5.4 Regional Dominance

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 2.11 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.[45]

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.

[45]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 2.11: Regional Robustness

	Annual employment growth		
	(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, %		
	(1)	(2)	(3)
	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 statisitics.

2.5.5 Firm Dynamics

Figure 2.11 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 2.12 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 2.11: Establishment-Level Employment Change (2009 - 2010)

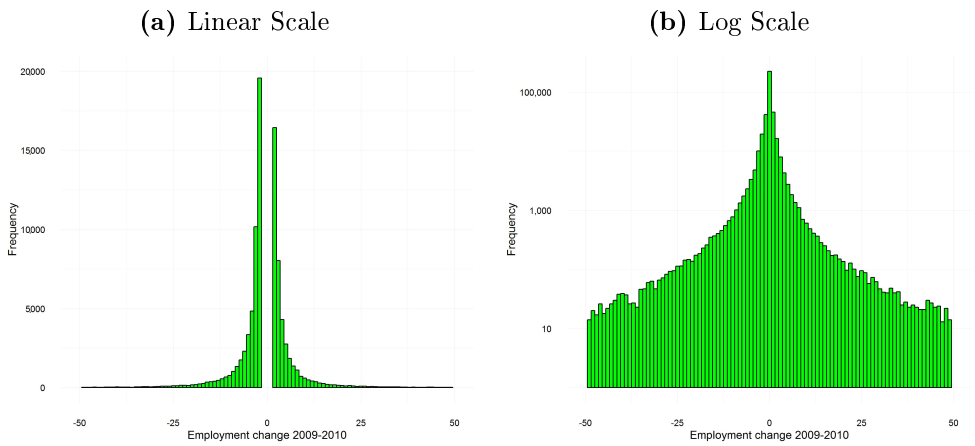

Source: ONS. Locally non-tradable firms only. Bars representing fewer than ten establishments have been excluded.

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 the nodal employment of establishments that existed in both years - the intensive margin. This is consistent with studies such as Giroud and Mueller (2017b) that use establishment-level regressions and log changes (since log changes at the extensive margin are undefined).

Figure 2.12: Total Employment Change (2009 - 2010)

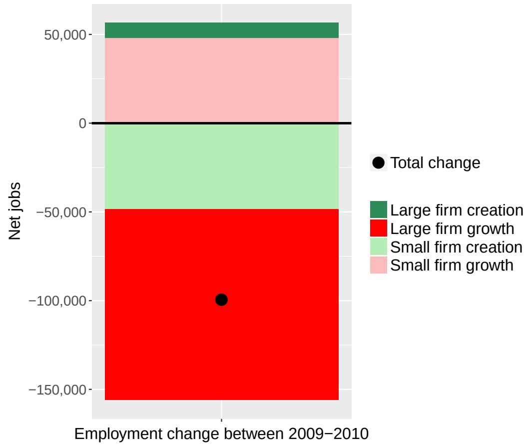

Source: ONS. This figure shows a decomposition of the total change in employment in the locally non-tradable sector between 2009 and 2010. Large firms are defined here as those with more than ten employees. Firm growth occurs when the establishment employs a non-zero number of people in both years. Firm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

Table 2.12 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 2.12, this suggests that we can attribute between half and two thirds of the cash-flow effect to firms making marginal changes to employment.

Table 2.12: Intensive Margin of Adjustment

	Annual employment growth		
	(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.

2.6 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, I claim it supported consumption, 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 chapter suggests that a 1pp fall in Bank rate was associated with around a 3.5pp increase in employment growth of locally non-tradable firms in the year after the shock. 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, which is around 700,000 jobs. This is a substantial estimated effect and comes with large caveats. Applying cross-sectional identification to an aggregate effect requires some strong assumptions, many of which are laid out in Nakamura and Steinsson (2017). To get a sense of how large this macroeconomic effect is we can compare it 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) develop a new measure of monetary shocks for the UK that purges interest rate changes of responses to economic fundamentals within the policy-maker's information set. The key result from that work is that a 1pp transitory monetary policy shock in a VAR framework 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.[46] My results are substantially larger, perhaps four times as large, but in May 2009 only 7% of people thought interest rates were likely to *rise a lot* over the next twelve months.[47] Perception that the cash-flow shock was permanent is consistent with a substantially larger effect.[48] Nevertheless, the overall takeaway from this calculation should be that cash-flow effects were important during the Crisis even if the central estimates might be inappropriate for calibrating small changes in monetary policy on aggregate variables. But given the asymmetry of consumption responses to positive and negative shocks, the effect I estimate could prove to be

[46]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.

[47]Question 6 of the Bank of England's *Public Attitudes to Inflation* survey.

[48]See Christelis, Georgarakos, Jappelli, Pistaferri, and Van Rooij (2017) for experimental evidence regarding consumption responses for transitory and permanent income shocks.

a lower bound in the face of an interest rate tightening cycle (e.g., Bunn, Roux, Reinold, and Surico (2018)).

The real contribution of this chapter then is to show how monetary policy can lead to heterogeneous employment effects across space, as well as time. Regions that 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. In many cases that helped shield the more vulnerable parts of the country from the adverse effects of the Great Recession. For example, in Cornwall, where there is a relatively large share of employment in locally non-tradable businesses that rely on the local community. 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. Although the duration of the stock is likely to adjust more slowly, smaller direct cash-flow effects in the face of rising interest rates might offset some of the asymmetry in consumption response.

Chapter 3

Local Demand Shocks and Firm Networks

I construct the establishment network of firms in the United Kingdom using micro-level data. I show that locally non-tradable firms increase employment in response to contemporaneous increases in house prices, and find evidence that these firms also respond positively to local demand shocks in other parts of the country. But for the majority of the past decade these responses are likely to be an artifact of moderate spatial demand spillovers rather than evidence of significant balance sheet or informational frictions. My results suggest that local demand shocks are likely to have limited macroeconomic consequences for overall employment.

3.1 Introduction

The majority of private-sector employment is accounted for by firms that have multiple establishments, yet little is known about how important intra-firm networks might be for the transmission of shocks across the economy. I fill this gap in the existing literature by combining spatially accurate business employment data with local-authority-level house price data to investigate the spatial propagation of local demand shocks.

Intra-firm establishment networks are most important for businesses that rely on nearby demand for their goods and services.[1] I therefore build the network of UK firms' geographic linkages and split firms into those that are in locally non-tradable and locally tradable industries.[2] Firms in locally tradable industries can be characterised as having a high degree of geographic concentration and a revenue stream that is independent of location of operation. Locally non-tradable firms exhibit less geographical concentration, primarily belong to the retail sector and rely on demand from people living close by.

I show that positive house price changes are associated with an increase in establishment-level employment for locally non-tradable firms. This is consistent with house price changes capturing information about local demand shocks, but I remain relatively agnostic about the precise mechanism that feeds through to employment changes. Following Giroud and Mueller (2017b), house prices proxy for local demand shocks in this chapter. At a minimum, I argue there is no reverse

[1] See Chapter 4, which shows locally non-tradable firms expand employment by opening up more establishment locations.

[2] For more details on firm classification, see Chapter 1 and Appendix A.

causation between employment shocks and house price growth because locally non-tradable employment makes up a relatively small fraction of the labour force.

I go on to show that positive local demand shocks spill over to nearby local authorities in ways that lead to positive employment effects across geographic regions. But I find little evidence that establishment networks matter for intensive-margin employment decisions through balance sheet or informational spillovers, or for locally tradable firms in general.

Large falls in house prices during the Great Recession in the US were associated with falls in consumer demand and led to employment losses in the locally non-tradable sector (e.g., Mian and Sufi (2014) and Giroud and Mueller (2017b)). These case studies are helpful to aid our understanding about firm behaviour in the Crisis, but are less informative about employment responses by these firms in the past decade. In my central specification I use OLS regressions with time and firm fixed effects for my establishment-level employment data between 2010 and 2017. I find that the that employment semi-elasticity for locally non-tradable firms with respect to contemporaneous house price changes is around 0.05 percentage points (pp), slightly below estimates in the US during the Crisis (e.g., Giroud and Mueller (2017b)).

I go on to test whether local demand shocks in other parts of the country affect local employment decisions. For example, does a house price shock in Reading have any impact on the employment decision of a restaurant in Oxford? I find evidence that establishments belonging to large and well-connected firms are particularly affected by demand shocks in other parts of the country. Specifically, I estimate that an increase in house prices in a separate local authority where the firm has a sister establishment is associated with an employment effect of up to a quarter of the local demand effect, and in the same direction.

This elasticity is difficult to interpret but could capture collateral constraints, information flows or customer-base measurement issues. By analysing different subsets of the data, I show that the most likely explanation for this apparent network effect is a spatial demand spillover caused by people moving across local authority boundaries. For example, there is no evidence that a house price shock in Glasgow has any effect on the employment decisions at the restaurant in Oxford, though a house price shock in Reading might. My results survive a range of robustness tests including using local-authority aggregated employment data, using lags of house price growth as instruments following Arellano and Bond (1991), and using council tax rates to proxy for local demand shocks.

One recent area of attention in research on firm networks has been the analysis of how shocks to input-output networks might reverberate around the system. Acemoglu, Akcigit, and Kerr (2016), Bigio and La'O (2016) and Baqaee and Farhi (2017) demonstrate the importance of firm supply chains for aggregate output, while Baqaee (2016) examines the importance of input-output networks at the extensive margin. One important constraint facing empirical studies of input-output networks is the quality of data available. While sectoral data exist for developed countries, firm-to-firm linkages are currently rarely collected and therefore make a granular study challenging. I therefore move away from supply-chain networks and exploit the available UK data by looking at establishment-level employment networks instead.

In general there is better data availability for firms' regional networks, where this chapter fits in naturally. Studies such as Blanchard and Katz (1992) and Fogli, Hill, and Perri (2012) have examined the role of economic geography in the propagation of shocks and even Tahbaz-Salehi, Carvalho, Nirei, and Saito (2017) contains a flavour of geographic networks when looking at supply chain disruptions following the 2011 Japanese Earthquake. Tahbaz-Salehi, Carvalho, Nirei, and Saito (2017) use the fact that physical production is not homogeneous across the US by modelling the cost of economic transactions as increasing in distance. Many of these studies focus on the role of productivity shocks for aggregate outcomes, either at the regional or sectoral level. But finding a reliable source of regional variation in total factor productivity is empirically challenging. This chapter takes an economic geography approach to understanding firm networks but focusses on the demand side. Instead of using shocks to total factor productivity or variation in supply-chain frictions, I use variation in regional demand to investigate firm employment behaviour.

A large body of work has looked at local demand shocks in the US, primarily by exploiting the cross-sectional variation in house prices. For example, Mian, Rao, and Sufi (2013) and Kaplan, Mitman, and Violante (2016) demonstrate the link between household balance sheets and consumption and go on to show the effect this link had on employment in industries that relied on local custom in Mian and Sufi (2014). Other studies such as Stroebel and Vavra (2014) have reaffirmed the importance of housing as a determinant of local economic variables. But less work has been done on the link between local demand and house prices following the Crisis, where this chapter makes a contribution.

The importance of the liability side of UK household balance sheets and the link between collateral and consumer behaviour is well-established (e.g. Bunn and Rostom (2015) and Aron, Duca, Muellbauer, Murata, and Murphy (2012)). On the firm side, there is some evidence that UK house prices explain a sizable fraction of fluctuations in corporate credit conditions and unemployment rates as Pinter (2015) shows, perhaps through a residential collateral mechanism affecting firms via company directors in Bahaj, Foulis, and Pinter (2017). But these studies are less suited to separating out how local demand shocks might propagate between or within firms.

The central aim of this chapter is to further our understanding regarding the spatial dimension of local demand shocks and investigate whether shock transmission is driven by spatial or organisational linkages. One possible reason for the latter being important would be if firm credit constraints mattered at the establishment level. Giroud and Mueller (2017a) show that highly levered firms cut back more on employment in the Great Recession but this chapter is closest in spirit and methodology to related work in Giroud and Mueller (2017b), which looks at how local demand shocks propagated through firms' internal networks in the Great Recession.

The rest of this chapter is organised as follows. Section 2 dives straight into an empirical analysis of how local demand shocks affect locally non-tradable employment decisions. Section 3 analyses how these shocks might propagate within firms and across space. I perform some robustness checks of my results in Section 4, and Section 5 concludes. For more information on the data sources please see Chapter 1.

3.2 Local Demand Shocks

Changes in house prices potentially lead to wealth effects for home owners and collateral effects for mortgagors. Since rents move more slowly than house prices, the net effect of a positive house price shock is likely to be positive for consumption.[3] Previous work for the US has shown that house price falls were an important factor weighing down on spending during the Great Recession and areas that suffered house price falls experienced falls in locally non-tradable employment, which serves as a proxy for local spending when consumption data are unavailable at a granular level (e.g., Mian and Sufi (2011); Mian and Sufi (2014); and Giroud and Mueller (2017b)).

In my analysis I use annual house price changes across UK local authorities to exploit regional variation. Unlike the studies mentioned above, I also employ a panel approach and use fixed effects to account for unobservable heterogeneity in the estimated effects at the local-authority and firm level. Spanning more years also helps to broaden my conclusions away from the Crisis period. Since I look at establishment-level employment changes it is advantageous to use a growth rate

[3] Though there is substantial micro evidence that collateral effects are likely to be much stronger than wealth effects (e.g., Aron, Duca, Muellbauer, Murata, and Murphy (2012)). To the extent inflated asset values might necessitate larger deposits, consumption might in fact fall in response to positive house price shocks.

that is robust to large negative percentage changes throughout.[4] As in Chapter 2, I therefore define locally non-tradable employment growth ($e_{i,t}$) and house price changes ($hp_{i,t}$) as follows

$$e_{i,t} = 2 \times \frac{E_{i,t}^{NT} - E_{i,t-1}^{NT}}{E_{i,t}^{NT} + E_{i,t-1}^{NT}} \tag{3.1}$$

$$hp_{i,t} = 2 \times \frac{HP_{i,t} - HP_{i,t-1}}{HP_{i,t} + HP_{i,t-1}} \tag{3.2}$$

Giroud and Mueller (2017b) regress the change in house prices between 2006 and 2009 on the change in employment at locally non-tradable establishments between 2007 and 2009 and find an employment semi-elasticity of a little over 0.1. So they estimate that a 10% fall in house prices was associated with a fall in locally non-tradable employment growth of around 1pp.

[4] Log differences become an increasingly poor approximation as the percentage change approaches -100%. The following growth rates are similar to log change for over small intervals but are much more accurate for large percentage falls (as is often the case for establishment-level employment). The growth rate definition in Equation 3.1 means that I can also calculate growth rates for ELAs that disappear entirely by setting their employment to zero in the second year. The alternative, which is to drop ELAs that do not appear in both years, yields qualitatively similar results. Using log changes would restrict attention entirely to the intensive margin (since non-zero employment is required in both years).

3.2.1 Establishment-Level Local Demand Effects

My establishment-level data spans 2008 to 2017 and my initial specification is shown below in Equation 3.4.

$$e_{x,i,t} = \eta_1 hp_{i,t} + \alpha_x + \delta_t + \varepsilon_{x,i,t} \tag{3.3}$$

Annual employment growth, $e_{x,i,t}$, at establishments belonging to firm x in local authority i at time t is explained by contemporaneous annual house price growth, $hp_{i,t}$, in the same local authority.[5] I focus attention on firms that operate in at least two local authorities but show the results for other firms in the robustness checks. Some multi-local-authority firms have more than one establishments in each local authority. To simplify the analysis I sum up firm employment by local authority and create aggregated enterprise-local-authority (ELA) units of observation and so $\{x,i\}$ corresponds to a particular ELA.[6] In the following regressions, all specifications are weighted by ELA employment size and double clustered at the local authority and enterprise level following standard practice. The reported number of observations corresponds to the number of ELA-year data points in my data. Although there are many more small firms, large firms have many more establishments, which means the two samples often have a similar number of total observations.

[5] I experimented with a series of lag structures. The ones presented gave the most convincingly intuitive results.

[6] While this makes the analysis slightly simpler, it may be true that one large supermarket has different employment characteristics to two city-centre supermarkets operating in the same local authority. Nevertheless, at this stage is does not have any bearing on the results because house price growth is measured at the local-authority level.

I include firm fixed effects, α_x, and time fixed effects, δ_t . The use of establishment data means that the employment growth rates better capture the intensive margin of adjustment (firms employing more staff) rather than the extensive margin (firms opening up new establishments).

The main parameter of interest in Equation 3.3 is η_1, which represents the semi-elasticity of employment with respect to a change in house prices in the average local authority.[7] The first column of Table 3.1 shows that the η_1 is around 0.05, which is around half of the non-panel estimate of Giroud and Mueller (2017b).

The second and third columns of Table 3.1 split the sample into local authorities that experience negative and positive house price growth rates. These suggest that house price growth might have an asymmetric effect on locally non-tradable employment. More specifically, the second column takes ELA-year observations for when house price growth were negative. The estimated coefficient suggests that a 10% fall in house prices was associated with around a 1.8pp fall in locally non-tradable employment growth and this effect is significant at the 1-percent level. This stands in contrast to the third column that suggests a 10% rise in house prices was associated with only a 0.6pp increase in locally non-tradable employment growth. Although the coefficient in insignificantly different from zero, the magnitude is still plausible.

Although I remain agnostic on the precise mechanism at work, the asymmetric effect of house price growth implied by these results is consistent with collateral shocks being important. Home equity withdrawals during the Great Recession

[7] It is a semi-elasticity because the outcome variable measures percentage point changes in employment growth rather than percentage changes in employment.

Table 3.1: Establishment-Level Local Demand Effects

	Locally non-tradable employment annual growth		
	(1)	(2)	(3)
	All data	-ve HP growth	+ve HP growth
Local HP growth, %	0.045*	0.184***	0.058
	(0.033)	(0.063)	(0.048)
Years	08-17	08-17	08-17
Fixed effects	Yes	Yes	Yes
Specification	OLS	OLS	OLS
Observations	610,258	148,613	424,930
R-squared	0.119	0.325	0.180
Note:		*p<0.1; **p<0.05; ***p<0.01	

All regressions are weighted by employment and standard errors (in parentheses) are double clustered at the firm and local-authority level. Fixed effects includes time and local-authority effects.

collapsed in the UK, which is consistent with the hypothesis that impediments to (cheaper) refinancing were a shock to disposable income and hence consumption choices. In turn, these fed in to locally non-tradable firms hiring fewer staff. Small rises in house prices since 2010, combined with a relatively subdued refinancing market, may have done little to slacken financial constraints and boost consumption. This might be thought of as a house-price-effect analogue of the results in Bunn, Roux, Reinold, and Surico (2018).[8]

3.2.2 Local-Authority-Level Local Demand Effects

It is useful to investigate whether the local demand effect and its asymmetry found above is also present at the local-authority level. My next specification regresses the percentage change in locally non-tradable employment in local authority i on the percentage change in house prices in the same local authority over the

[8]The effect of local demand shocks may be non-linear as well as asymmetric and the sample containing house price falls contains some very large negative growth rates. Alternative specifications of Equation 3.3 including squared terms support the claim the effect might be asymmetric.

past year, as shown in Equation 3.4. The key difference between Equation 3.3 and Equation 3.4 is that the latter is much better able to capture the extensive margin of adjustment.

$$e_{i,t} = \alpha + \eta_1 hp_{i,t} + \beta_i + \delta_t + \varepsilon_{i,t} \tag{3.4}$$

γ_i are local-authority fixed effects that capture time-invariant features of local authorities that might affect locally non-tradable employment dynamics and δ_t are year fixed effects that capture common shocks across time. Standard errors are clustered by local authority. Following Giroud and Mueller (2017b), I can include industry shares as controls but it has little effect on the results.

Table 3.2 shows the results when I look at area-wide locally non-tradable employment, without restricting attention to surviving establishments. Each year-local-authority observation is simply the sum of employment at locally non-tradable establishments in that area so employment growth rates capture the intensive and extensive margin of adjustment.

The first column uses the complete panel and applies time and local authority fixed effects. The main coefficient of 0.057 is reasonably similar to the estimate for η_1 at the establishment level and is now significant at the 5-percent level. Running the same regression without local authority fixed effects yields an estimated coefficient of 0.097 with a high degree of precision but a Hausman tests confirms that the fixed-effect estimator is preferred to the random effects estimator because at least the former is consistent, even if it might be inefficient.

Table 3.2: Local-Authority Level Local Demand Effects

	Locally non-tradable firms, annual growth			
	(1)	(2)	(3)	(4)
	Employment	Employment	Employment	Establishments
	All data	-ve HP growth	+ve HP growth	All data
HP growth, %	0.057**	0.095	0.076	0.005
	(0.029)	(0.077)	(0.051)	(0.016)
Years	08-17	08-17	08-17	08-17
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Observations	3,890	1,155	2,735	3,890
R-squared	0.061	0.046	0.069	0.233
Note:			*p<0.1; **p<0.05; ***p<0.01	

Standard errors (in parentheses) are the local-authority level. Fixed effects includes time and local-authority effects.

The second and third column once again split the sample into local authorities that experience negative and positive house price growth. When looking at local authorities as a whole, both coefficients are not statistically significantly different from zero (and from one another) but have plausible magnitudes. This suggests that the local-authority-wide response of locally non-tradable employment to house prices might not be asymmetric around zero house price growth, contrary to the findings in Table 3.1. One interpretation of this is that firms cut back on hiring during a crisis but the negative local demand shock does not lead to an asymmetrically large effect on the number of establishments that close down. In other words, the intensive margin of adjustment appears to behave asymmetrically with regards to the direction of house prices.

The final column repeats the analysis in Equation 3.4 but replaces the dependent variable with the growth rate of the number of establishments in each local authority. The purpose of this test is to see if employment or the number of establishments is the more interesting avenue to investigate. The fixed effects estimator shows there is no discernible relationship between house price changes and the net number of firms born every year since the estimate is both small and statistically insignificantly different from zero. Once again, the alternative semi-elasticity estimate only including time fixed effects was large and significant but rejected by a Hausman test. The remainder of the analysis focuses on employment because it is a more flexible margin of adjustment and there is a clearer mapping from demand shocks to firm decision making.

Overall, the takeaway from this section is that local demand effects, as captured by changes in house prices, are slightly shy of those estimated by Giroud and Mueller (2017b) in the US but of the right order of magnitude. There is some evidence that the effect might be larger in magnitude at the intensive margin for house price falls so it is worth allowing for this possibility when investigating spatial propagation. It is, however, difficult to capture asymmetric house price shocks when we move to looking at firm network effects so from hereon I look at the Crisis period (2008-2009) and the post-Crisis period (2010-2017) separately as a reasonable alternative.

3.3 Spatial Propagation

In this section I move beyond the effect of local demand shocks and explore whether the establishment structure of locally non-tradable firms can tell us anything about how shocks propagate across the economy. The first part briefly summarises some interesting features of UK spatial linkages before I go on to test how demand shocks might propagate across firm establishment networks.

3.3.1 Geographic Linkages

The average local authority contains 6,500 ELAs, although this varies widely across the country. At one end of the spectrum, the London Borough of Westminster, Birmingham and Leeds all contain more than 50,000 ELAs. At the other, many Scottish islands and other extremities of the UK contain fewer than 100 ELAs. The vast majority of people employed by locally non-tradable firms in the major conurbations work for enterprises that operate in multiple authorities around the UK. Many other multi-local-authority enterprises have a small number of establishments in contiguous authorities. A natural example might be a car mechanic who expands her business by opening up another garage in the next town along, which happens to be in a contiguous local authority. The average local authority in the UK has five direct neighbours and so it may not be surprising that the average firm operating in multiple local authorities has around four ELAs as part of its network. Conditional on operating in more than two authorities, the average firm employs a little under 40 people in each.

Taking local authorities as a whole, firm establishment networks in the UK are pervasive. Given the definitions of locally non-tradable and tradable firms in Chapter 1 and Appendix A, it is unsurprising that locally non-tradable firms tend to have more expansive geographic networks. In fact, there are around ten times more locally non-tradable ELAs than locally tradable ELAs. All local authorities contain firms that employ at least ten people in every one of the 389 local authorities; this is referred to as the *Starbucks effect* by Giroud and Mueller (2017b) (also appropriate terminology in the UK) and reflects the large number of retail chains and other locally non-tradable brands operating throughout the country. Even Orkney boasts a Tesco, Lidl, Boots and a Co-op.

Excluding locally non-tradable firms operating in a single local authority, the average locally non-tradable firm is connected to just under a third of all local authorities. Within this, there is substantial regional variation and Figure 3.1a shows that a high degree of connection exist in the Midlands, the area to the west of Bristol and parts of Scotland to the south of Glasgow. The high average number of connections for the Scottish islands reflect the relatively limited number of locally non-tradable firms (but those that are there have extensive linkages), though it stands in contrast to the other extremity of Cornwall. Northern Ireland stands out as being remarkably insulated from the rest of Great Britain and central London is surprisingly unconnected, perhaps reflecting the mix of small and large businesses.

The second panel, Figure 3.1b, shows the average number of local authority connections where the firm employs more than 100 people. This map is similar to Figure 3.1a with the notable exception that Aberdeenshire has quite a few connections to large firms even though it has fewer total connections on average. In an absolute sense, local authorities are well connected: the average number of regional connections for firms in even the least connected local authorities is above 40.

Figure 3.1: Geographic Networks

(a) Avg. Number of Connections

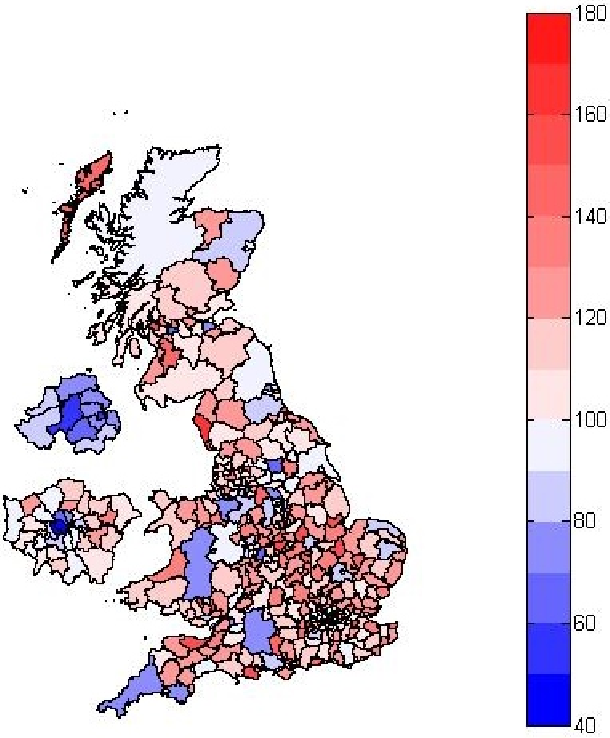

(b) Avg. Connections >100 Employees

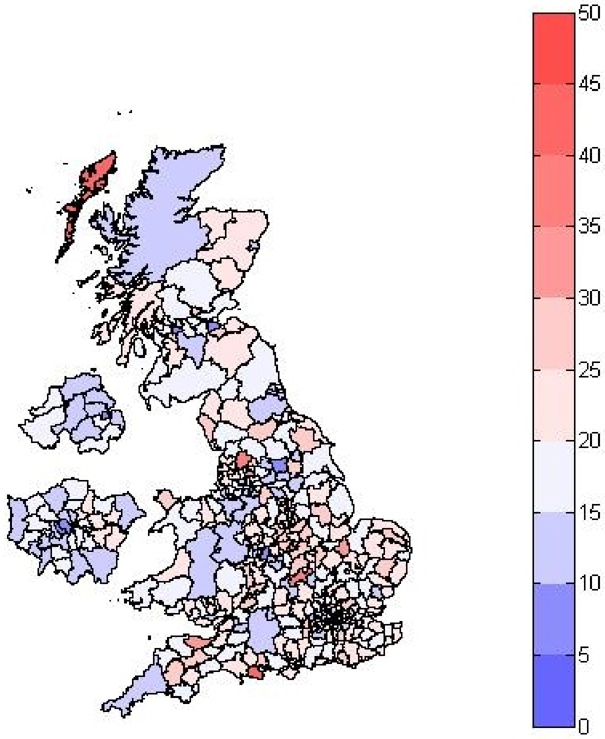

Source: ONS. These maps give an indication of how connected parts of the country are based on their locally non-tradable employment. The first map includes all locally non-tradable firms and the second only those that employ more than 100 people across their establishment network. Greater London is shown inset below Northern Ireland.

3.3.2 Establishment-Level Network Effects

Economic intuition would suggest that locally non-tradable establishments are likely to make employment decisions primarily based on local demand factors where possible. Demand shocks in other parts of the firm's regional network are, at first glance, not obviously relevant for that establishment's allocation of resources.[9] The next specification tests whether this is true by adding a *network term* to Equation 3.4, which captures the establishment's exposure to local demand shocks in other parts of its regional network. To do this I construct the establishment network of every locally non-tradable firm in the UK. As above, I aggregate the employment of each firm by local authority so each firm has a maximum network of 389 ELAs. For firm x in local authority i, their local demand shock is the growth rate of house prices in the immediate local authority as before. I go on to define their network demand shock as the weighted average annual house price growth in their establishment network, lagged by one year to allow for pass-through.[10] An example is shown in Figure 3.2.

Following on from Bailey, Cao, Kuchler, and Stroebel (2016) in the context of social networks and Giroud and Mueller (2017b) in the more direct context of identifying non-local house price demand effects, the initial specification in this part uses the simple linkage weights based on the the proportion of employment

[9]Nevertheless, strategic decisions to shut entire stores might be more likely to be taken by head office.

[10]I tried a number of specifications for the lag structure of Equations 3.3 and 3.6 and the results are reasonably sensitive to them. Under the hypothesis that local demand shocks transmit across firm establishment networks due to credit or informational frictions, this specification with a lagged network term seems reasonable.

Figure 3.2: Demand Shock Example

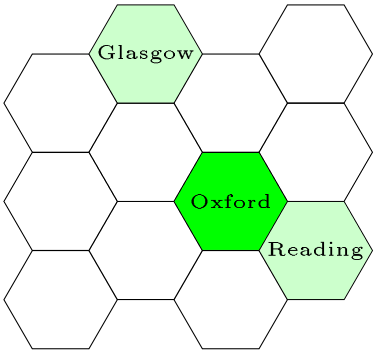

A restaurant chain has three establishments in the UK. For the Oxford branch, the local demand shock is defined as the change in house prices in local authority of Oxford (dark green). The network demand shock is defined as the weighted average change in house prices in the other two local authorities of Reading and Glasgow (light green). Those weights depend on the relative employment between the other two establishments.

in other local authorities.[11] For firm x in local authority i, each distant local authority j has a weight $\omega_{x,i,j,t}$ at time t based on the relationship in Equation 3.5.[12]

$$\omega_{x,i,j,t} = \frac{Emp_{x,j,t}}{\sum\limits_{j \neq i} Emp_{x,j,t}} \tag{3.5}$$

This means that a firm employing the same number of people in three local authorities would put an equal weighting on the other two region's house price changes when calculating the weighted average network term. In Equation 3.6, the coefficient η_2 captures the elasticity of employment with respect to a weighted average house price growth in the other local authorities the firm operates in. The null hypothesis is that $\eta_2 = 0$ and the network-weighted change in house prices has no effect on employment decisions.

[11] The following analysis is also similar in spirit to Dupor and Guerrero (2017) in the context of government spending shocks.

[12] The results are qualitatively similar if I define the time-invariant weights at the start of my sample. But it does seem appropriate to update weights as firm structure evolves.

$$e_{x,i,t} = \eta_1 hp_{i,t} + \eta_2 \sum_{j \neq i}^{J} \omega_{x,i,j,t} hp_{i,t-1} + \alpha_x + \delta_t + \varepsilon_{x,i,t} \tag{3.6}$$

Since Section 3.2 suggested that demand effects might be asymmetric, I split the sample into two dimensions. As before, I define the Crisis period as 2008 and 2009 and the post-Crisis period as 2010 to 2017. Within each period I look at small firms and large firms, where the latter is defined as those with more than 1,000 employees across the entire network.[13] Understanding the differences between large and small firms is important because large firms are less likely to be credit constrained and surely better positioned to devote resources to understanding macroeconomic developments relevant for their business.

It is worth emphasising that these initial results focus again on the intensive margin, which is likely to be much more important for large firms since they have a lower probability of failure (and few firms are born large). In particular, the fixed costs of hiring and firing are likely to be greater for small firms so they should respond to demand shocks in different ways. Perhaps a small restaurant grants fewer shifts to its workers; delays hiring instead of actively firing; or tries to smooth out workflow with informal arrangements. Large firms might appear to adjust staffing changes in a smoothly continuous manner because of aggregation effects.

The interpretation of η_2 is not as straightforward as η_1 because a one percentage point increase in the network-weighted house price is not comparable to a

[13] 1,000 employees is roughly the median sized firm in this sub-sample. The median number of ELAs is around 250. Chapter 4 shows these large firms employ around a half of workers in the locally non-tradable sector.

one percentage point increase in the local house price: the latter is a much larger macroeconomic shock when networks are expansive. One useful way of interpreting η_2 is to consider the impact of a one percentage point increase in house prices in one of the networked local authorities. On average, large firms are connected to around 222 other local authorities so I divide the network term by this number.

The first column in Table 3.3 shows that there is little evidence in support of the existence of a network effect for small firms during the Crisis period. The second column shows there is weak evidence for a negative network term for large firms during the Crisis period. Taken at face value, this would imply that a negative demand shock at a distant, sister establishment would lead to a positive employment effect. That might be consistent with firms substituting employees across their establishment network but this is contrary to the findings of Giroud and Mueller (2017b). One candidate explanation for this in the UK would be the behaviour of firms in response to rent reviews. Between 2007 and 2009, increases in property values might have led to increases in rental costs. Property valuation benchmarks for business rates calculations also reset in 2010, to 2008 prices. Future expectations of large business rates bills might have led to intra-establishment employment substitution. Importantly, though, the coefficient is small and implies that a 10% fall in house prices in one local authority increases employment growth by a mere 0.02pp.[14]

The final two columns suggest that in the period after the Crisis large firms experienced a sizable positive network effect and small firms did not. For these

[14]The appropriate calculation is $-10 \times -0.509/222 \approx 0.02$pp.

large firms, a 10% house price increase was associated with a 0.06pp increase in employment growth, roughly a tenth of the local demand effect for local authorities as a whole. To put this coefficient in context, in 2010 the standard deviation of house price growth across local authorities was around 4.6%. That compares to the standard deviation of the network effect for large firms of around 0.8%, which is substantially lower because of aggregation effects.[15]

A similar pattern of results is shown in Table 3.4 for firm size based on number of connected ELAs rather than total employees. For well-connected firms after the Crisis, defined as those that have establishments in more than 250 local authorities, the network term represents around a quarter of the local-authority local demand effect estimated in the previous section.[16]

I proceed by focussing attention on the final column of Tables 3.3 and 3.4. Why do large and well-connected firms respond positively to local demand shocks in other parts of the country in non-crisis periods?

A positive network coefficient for large and well-connected firms could reflect one of three hypotheses. First, it could represent the propagation of shocks across firm networks due to firms being unable to allocate resources optimally in the presence of binding credit constraints, as suggested in Giroud and Mueller (2017b). For example, if a locally non-tradable firm suffered a severe negative demand shock in a sister, distant establishment, that might reduce its ability to pay its workers

[15] I divide the network coefficient by the average number of connections in the network as a form of scaling. To my knowledge, there is no accepted best practice when attempting to scale a network term such as this.

[16] That is to say, $4.528/313 = 0.0144 \approx 0.06/4$. It is closer to a third of the local demand effect estimated in Table 3.1. The local demand effect are often negative in the following regressions, though they are not precisely estimated. That could partly be due to multicolinearity resulting from the spatial correlation of house price changes across nearby local authorities.

Table 3.3: Establishment-Level Network Effects: Firm Size

	Locally non-tradable employment annual growth			
	(1)	(2)	(3)	(4)
	Crisis	Crisis	Post-Crisis	Post-Crisis
	Small firms	Large firms	Small firms	Large firms
Local HP growth, %	−0.001	−0.158	0.087**	−0.055
	(0.094)	(0.107)	(0.041)	(0.046)
Network HP growth, %	−0.100*	−0.509**	−0.048	1.361***
	(0.057)	(0.247)	(0.042)	(0.320)
Years	08-09	08-09	10-17	10-17
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Average ELA connections	26.8	222	28.0	222
Observations	48,774	57,908	235,991	252,177
R-squared	0.462	0.488	0.192	0.144
Note:				*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are double clustered at the firm and local-authority level. Fixed effects includes time and local-authority effect. Large firms are defined as having more than 1,000 employees.

Table 3.4: Establishment-Level Network Effects: Firm Connectedness

	Locally non-tradable employment annual growth			
	(1)	(2)	(3)	(4)
	Crisis	Crisis	Post-Crisis	Post-Crisis
	Small network	Large network	Small network	Large network
Local HP growth, %	−0.014	−0.243	0.057	−0.084
	(0.090)	(0.170)	(0.040)	(0.070)
Network HP growth, %	−0.199***	−1.497	0.113	4.528***
	(0.066)	(1.969)	(0.113)	(0.590)
Years	08-09	08-09	10-17	10-17
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Average ELA connections	78.9	310	73.9	313
Observations	81,306	24,386	378,793	110,854
R-squared	0.463	0.522	0.183	0.132
Note:				*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are double clustered at the firm and local-authority level. Fixed effects includes time and local-authority effects. Large firms are defined as having a network of more than 250 ELAs.

in the unaffected local authority. This would be especially likely if the firm was unable to raise financing or if they had an insufficient buffer of liquid assets held centrally. The problem with this narrative is that, outside the Crisis period, this mechanism is less convincing in reverse, especially for large firms that are much less likely to be close to their financing constraint. Although some large firms had trouble raising financing from after 2010, one would expect small firms to be closer to their financing constraint.

Another hypothesis for a positive network effect is that it might reflect managers exploiting the information contained within their firm network to inform their choices regarding hiring decisions in different locations. This could be true at the local-management level (in the case of franchised firms) or at head office.[17] In this scenario, firms use information from their branch network to elicit a picture of the aggregate or regional economy today and in the future. Under this scenario the network coefficient is therefore positive because other parts of the country provide information for the future prospects of a given establishment. The obvious example in this case is for firms close to London to use information on the London housing market to form expectations about potential ripple effects of house price growth across the country.

One final hypothesis is that the network term is an artifact of the local demand sphere of influence being inadequately defined by local authority boundaries. Put another way, house price changes exhibit regional patterns and so local demand shocks for establishments in nearby local authorities are likely to be spatially cor-

[17] Unfortunately the BSD employment data do not allow me to distinguish between franchised and non-franchised firms, let alone where management decisions are actually taken.

related. People also travel between nearby local authorities to purchase goods and services. This hypothesis is tantamount to the idea that establishments are primarily linked by realised spatial shocks rather than balance sheets or information flows.

One way of differentiating between these hypotheses is by focusing attention on a subset of large firms' establishment networks. In the first column of Table 3.5 I examine all large-firm ELAs within a 50km radius and set all other weights in Equation 3.6 to zero. In the second column I do the reverse and only capture the demand shocks in ELAs more than 50km away. The final two columns do the same thing for a cut-off of 100km.[18] The results suggest that the network effect for large firms only exists for ELAs that are reasonably close by, which is evidence in support of the importance of the spatial dimension rather than a story based on collateral frictions.

The results in Table 3.5, however, do not tell us whether employment decisions are made in expectation of future local economic conditions or as a response to realised ones. To test between the latter two hypotheses I therefore restrict attention to ELAs in the Home Counties in Table 3.6.[19] After 2008, London house prices were a leading indicator of regional house prices and there was a noticeable ripple effect of house price shocks spreading to the surrounding local authorities.[20] Moreover, very few people live in London and shop in locally non-

[18]All distances are measured from the centre of the local authority.

[19]For the purposes of this study I define the Home Counties as Surrey, Kent, Essex, Hertfordshire and Berkshire as these border Greater London. Greater London includes the ceremonial county of Middlesex.

[20]For example, see Hudson, Hudson, and Morley (2018) for a recent investigation. I include all years in years in the Home County regressions, partly to increase the sample size.

Table 3.5: Distance Effects

	Locally non-tradable employment annual growth			
	(1)	(2)	(3)	(4)
	Within 50km	Beyond 50km	Within 100km	Beyond 100km
Local HP growth, %	−0.049	−0.053	−0.046	−0.057
	(0.048)	(0.046)	(0.047)	(0.047)
Network HP growth, %	0.341***	0.877**	0.407***	0.285
	(0.072)	(0.348)	(0.075)	(0.288)
Years	10-17	10-17	10-17	10-17
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Observations	243,698	251,977	250,683	251,410
R-squared	0.144	0.143	0.143	0.142
Note:				*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are double clustered at the firm and local-authority level. Fixed effects includes time and local-authority effects. Results are shown for large firms, defined as having more than 1,000 employees. The network term is calculated by taking a weighted average of networked local authority house price changes as before and setting weights to zero when the centre of that local authority is either within or beyond the distance cut-off.

tradable firms in the Home Counties: the flow of people goes the other way for commuting purposes.[21]

The first two columns of Table 3.6 show that the usual pattern for small and large firms is true for this subsample of ELAs; the network term is only significant for the large firms. The third column replicates the analysis for the large firms but zeroes out the the weighting of local authorities in Greater London. The estimated magnitude of the network term falls but it is still statistically significant at the 1-percent level. The final column performs the opposite calculation and sets the weights of local authorities to zero unless they are in Greater London and the statistical significance of the network term disappears entirely.

The last three columns of Table 3.6 suggest that the network term might be

[21]There are likely to be a few exceptions such as out-of-town shopping centres.

Table 3.6: Establishment-Level Network Effects in the Home Counties

	Locally non-tradable employment annual growth			
	(1)	(2)	(3)	(4)
	Smal firms Full network	Large firms Full network	Large firms exc. London network	Large firms only inc. London network
Local HP growth, %	−0.033	−0.345	−0.346	−0.335
	(0.082)	(0.201)	(0.201)	(0.205)
Network HP growth, %	−0.070	1.583***	0.898***	0.130
	(0.150)	(0.331)	(0.262)	(0.251)
Years	08-17	08-17	08-17	08-17
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Observations	37,072	46,692	46,684	46,227
R-squared	0.182	0.110	0.109	0.109
Note:				*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are double clustered at the firm and local-authority level. Fixed effects includes time and local-authority effects. For the last two columns the network term is calculated by taking a weighted average of networked local authority house price changes as before and setting weights to zero if the local authority is beyond or within the Greater London area, respectively.

driven by demand spillovers rather than information transmission. The significant network term in the penultimate column suggests that there were local demand spillovers in the local authorities surrounding those establishments in the Home Counties. But the insignificant term in the last column suggests that, although firms were aware of house price shocks in London, they did not adjust employment as a result. That in turn suggests the network term probably reflects regional demand spillovers more than expectation formation across establishment networks. Put another way, the network term might still be capturing a local demand shock. So the most convincing reason I find that demand shocks matter is that people travel across administrative borders; demand shocks ripple across the countryside.

3.3.3 Macro Network Effects

In this part I redo the analysis for local authorities as a whole to investigate the implications for macroeconomic propagation. As stated above, this approach is better able to capture both the intensive and extensive margin of adjustment. To look at the aggregated network effects, I create an aggregate analogue to the spatial weighting matrix in Equation 3.5.

$$\omega_{i,j,t} = \frac{\sum\limits_{x} Emp_{x,j,i,t}}{\sum\limits_{j \neq i} \sum\limits_{x} Emp_{x,j,i,t}} \tag{3.7}$$

Here, $\omega_{i,j,t}$ is the entry for the spatial weighting matrix for local authority i that has connections in local authority j at time t. The numerator is the sum of employment across establishments x in local authority j that have sister establishments in local authority i at time t. The denominator is the sum of intra-firm, extra-local-authority employment. As before, $e_{i,t}$ is the percentage employment growth defined in Equation 3.1, now at the local-authority level. I can then test the local-authority version of Equation 3.6.

$$e_{i,t} = \gamma_1 hp_{i,t} + \gamma_2 \sum_{j \neq i}^{J} \omega_{i,j,t} hp_{i,t-1} + \alpha_x + \delta_t + \varepsilon_{i,t} \tag{3.8}$$

Once again, I test the null hypothesis of Equation 3.6 that $\gamma_2 = 0$ and that non-local demand shocks do not affect local authority locally non-tradable employment growth. The first column of Table 3.7 shows the central results using the OLS estimation strategy from before. The coefficient for local house price growth is similar to previous estimates, at a little below 0.06.[22]

In the first column of Table 3.7, using the local authority locally non-tradable employment spatial weighting matrix for the complete network, the network coefficient is relatively large. It is somewhere between the estimates obtained in the establishment-level regressions in the fourth columns of Tables 3.3 and 3.4. As before, interpreting the size of the coefficient is complicated by the fact that most local authorities are connected to all 389 and the standard deviation of the local-authority network effect is between 25 and 60 times smaller than the standard deviation of the local demand effects because of the law of large numbers. Using the same sense check as before, I normalise the network term by the 389 local authorities. This suggests the impact of a 10% house price increase in a distant local authority leads to around a 0.06pp increase in locally non-tradable employment growth in the local authority of interest, which is around a ninth of the local demand effect.[23]

The second column in Table 3.7 tests this by excluding firms that employ more than 1,000 people from the employment and weighting matrix calculations. The insignificant coefficient once again suggests that network effects are primarily pertinent for large firms even when taking account of the extensive margin adjustment of small firms. The third column only includes establishments in local authorities that are more than 100km away in the weighting matrix. In this case

[22] In these regressions I use the full span of data available to maximise the number of observations. Results are similar, though less precise, when focussing on the post-Crisis period.

[23] That is, $10 \times 2.459/389 \approx 0.06$pp. Replacing the denominator by the average number of local authority connections large firms are connected to in Table 3.3 increases the relative network effect to around a fifth of the local demand effect.

Table 3.7: Local-Authority-Level Network Effects

	Locally non-tradable employment annual growth		
	(1)	(2)	(3)
	All data	Small firms	Network beyond 100km
Local HP growth, %	0.058**	0.054*	0.059**
	(0.029)	(0.029)	(0.029)
Non-local HP growth, % Complete network	2.459**		
	(0.986)		
Non-local HP growth, % Small-firm network		0.468	
		(0.431)	
Non-local HP growth, % Distant network			−0.317
			(0.273)
Years	08-17	08-17	08-17
Fixed effects	Yes	Yes	Yes
Specification	OLS	OLS	OLS
R-squared	0.062	0.063	0.062
Observations	3,890	3,890	3,890
Note:			*p<0.1; **p<0.05; ***p<0.01

Fixed effects includes time and local-authority effects and standard errors are shown in parentheses.

the network term is unlikely to reflect geographic demand spillovers. Consistent with the establishment-level results, this coefficient is not statistically different from zero, which supports the claim that the apparent network effect is driven by demand shocks in nearby locations.

3.4 Robustness Checks

In this section I conduct a series of robustness checks to test the main results. In the first part I use lags of house price growth rates as instruments in a GMM environment. I go on to test my central results using house price changes using subsets of firms and an alternative weighting structure. Finally, I consider using council tax changes to proxy for local demand shocks.

3.4.1 Local Demand Shock Identification

The regression analysis used above assumes that house price growth perfectly captures local demand shocks. Significant coefficients, however, do not necessarily imply a causal relationship between the effect of house prices and locally non-tradable employment. In this study I follow Giroud and Mueller (2017b) and remain relatively agnostic about the precise mechanism by which house prices affect local demand. So I do not make a distinction between the effect of an exogenous increase in house prices (it is unclear what that would even mean) or a more general shock that affects the demand for housing among other things. What I do need to rule out, however, is that employment fluctuations in the locally non-tradable sector affect house prices.

The most common way of addressing this identification challenge was developed in the seminal work of Mian and Sufi (2011) and uses the housing supply elasticity as an instrument for house price changes. This approach is valid if the instrument is uncorrelated with coincident drivers of housing wealth and consumption (e.g. regional exposure to certain industries) in the local area. Housing supply estimates exist for the UK (e.g. Hilber and Vermeulen (2016)) but the general approach has been criticised by Davidoff (2016) who argues that supply constraints might be correlated with endogenous characteristics that make some places relatively more attractive to live. This applies to physical determinants of housing supply such as topography, but also to man-made zoning restrictions. Local governments never make decisions at random and it is plausible that areas where planning permission is denied are often associated with higher house prices (for example, on the green belt).

I therefore use the Arellano-Bond dynamic panel estimator (see Arellano and Bond (1991)) to instrument for house price changes as a robustness check. This is especially suited to panels with a large number of cross-sectional observations and a small number of time periods, as is the case in my sample.

The fixed effects estimator eliminates the time-invariant component of the error term before applying OLS to the model. But for this approach to be valid we require the following.

$$hp_{i,t} \perp \varepsilon_{i,t} \, \forall t \in [1, T] \tag{3.9}$$

Intuitively, this requires that changes in locally non-tradable employment at no point feed back into changes in house prices. This is a relatively standard assumption in studies of this sort and can be justified by the relatively small fraction of total employment that locally non-tradable firms make up. The average local authority has between a fifth and a quarter of people employed in locally non-tradable firms. My key identifying assumption is therefore that exogenous shocks to the locally non-tradable sector have a very minor impact on the house prices in any local authority. It is, however, helpful to check this specification's validity using General Method of Moments (GMM) techniques.

Since house prices are persistent, it is important to take first differences in the annual growth rates. An Arellano-Bond test for second-order autocorrelation yields a z-value 0.66 (and a p-value of 0.51) so we cannot reject the hypothesis that the errors are not serially correlated in the differenced model. The test for first order autocorrelation is mechanically negatively significant because of the following relationship

$$\Delta e_{i,t} = \eta_1 \Delta hp_{i,t} + \Delta\gamma_i + \Delta\varepsilon_{i,t} \tag{3.10}$$

The efficient estimator can then be obtained by computing the two-step first-difference GMM estimator using lags and leads of $\Delta hp_{i,t}$ as instruments. For example, lagged house price growth will likely be a good instrument because it will be correlated with $\Delta hp_{i,t}$ but not with $\Delta\varepsilon_{i,t}$. Table 3.8 shows the results for a number of different specifications. The first column is my preferred specification. It uses all lagged values of house price growth rates as instruments and yields a statistically significant coefficient of 0.08, which is only slightly higher than that found in the simple fixed-effects regression. The Hansen over-identification test checks whether the instruments are valid and the null hypothesis is not rejected for either test in the first column, suggesting the instruments can be treated as exogenous.[24]

The second column shows what happens when only two lags of house price growth are included as instruments. The estimate of η_1 is lower and insignificant, which is unsurprising given we would expect more accurate results when more instruments are used. Furthermore, it would seem strange to think that distant lags

[24]In all regressions, the orthogonal deviations transformations described in Arellano and Bover (1995) leaves the coefficients unchanged.

of house price growth would be more likely to be linked to employment growth, so they may as well be included. The third column shows the results using a one-step estimator, which is only efficient when standard errors are homoskedastic. Instead of just using first differences, the fourth column uses the system-GMM estimator (see Arellano and Bover (1995)), which has the advantage that it allows for more instruments (and so can increase efficiency) at the expense of computational complexity. The coefficient remains approximately stable but the Hansen test suggests the instruments may be invalid. The fifth column uses all possible leads and lags as instruments, which requires error terms to be uncorrelated with all past and future house price changes. Although in principle using more instruments increases the accuracy of the estimates, the Hansen test suggests the full set might not be valid.

The key takeaway from the instrumental variables analysis is that, under the assumption that locally non-tradable employment is a small enough fraction of overall employment, all reasonable specifications using GMM estimates yield coefficients similar to those obtained in Table 3.2. This should give us confidence that using changes in house price to proxy for local demand shocks in the UK is a valid econometric approach.

3.4.2 Firm Types and Weighting

Section 3.3 uses multi-unit locally non-tradable establishment employment data to show that local demand shocks are not captured by local authority boundaries and regional demand spillovers lead to a positive network effect for nearby estab-

Table 3.8: Instrumental Variables

	Locally non-tradable employment annual growth				
	(1)	(2)	(3)	(4)	(5)
HP growth, %	0.080***	0.041	0.069**	0.067***	0.061**
	(0.030)	(0.031)	(0.032)	(0.026)	(0.030)
Years	08-17	08-17	08-17	08-17	08-17
Specification	GMM	Orthogonal	One-step	System	GMM
Instruments	All Lags	2 Lags	All Lags	All Lags	All Lags and leads
Hansen p-value	0.121	0.484	0.121	0.052*	0.060*
Diff. Hansen p-value	0.184	0.879	0.184	0.979	0.086*
Observations	3,501	3,501	3,501	3,890	3,501
Note:					*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and heteroskedastic-robust standard errors are shown in parentheses.

lishments. In this part I perform a few more specifications to test the robustness of this finding.

The first check is to examine locally non-tradable firms that are only present in one local authority. These single-ELA firms were dropped from my main regressions as they have no network coefficient by construction, but for completeness it is helpful to test whether local demand shocks lead to changes in employment. In fact, the coefficient in the (unreported) local demand coefficient is insignificantly different from zero with a t-value of 0.37.[25] In part, this is because these firms are very small and the intensive margin of adjustment is less likely to be used. To the extent local demand shocks operate on the extensive margin for these firms, that is captured in my macroeconomic-effects analysis above. But on the intensive margin it is reassuring to know that excluding these firms does not miss a key part of the local demand story.

Table 3.9 runs the standard regression on firms in the locally tradable sector, which serves as a placebo for my analysis. Local demand shocks proxied by

[25] Full robustness checks are available upon request.

house price changes affect locally non-tradable firms as they generate revenues from people who live close by. Locally tradable firms on the other hand might employ people locally but the demand for their goods and services should be entirely independent of their geographical location. We would therefore expect to find no local or network demand effects for locally tradable firms. The first two columns in Table 3.9 split firms into small and large firms across all time periods using the previous definition of 1,000 employees for the cut-off. The final two columns look at large firms during and after the Crisis. The insignificant local and network coefficients for all columns provides evidence that these firms do not adjust employment in the face of house price shocks and therefore the house price changes are indeed picking up local demand shocks.

Table 3.10 tests the importance of the weighting structure in Equation 3.5. In particular, I assign an equal weighting to all local authorities where the firm of interest has any presence, rather than weighting by employment exposure. Using a similar distance cut-off as before, I show that the network term is only significant when nearby local authority demand shocks are captured. Giroud and Mueller (2017b) found that their network coefficients disappeared when an equal weighting structure was used and they interpreted this as a successful placebo test for their collateral constraints hypothesis. In my case, the pattern of results presented in Table 3.10 is consistent with the spatial demand spillover hypothesis, though the equal weighting does lead to smaller coefficients. One interpretation is that the equal weighting is most appropriate for the spatial dimension. After all, the original weighting was constructed under the assumption that the sister es-

Table 3.9: Locally Tradable Firms

	Locally tradable employment annual growth			
	(1)	(2)	(3)	(4)
	All years Small firms	All years Large firms	Crisis Large firms	Post-Crisis Large firms
Local HP growth, %	0.090	0.027	−0.095	0.093
	(0.081)	(0.117)	(0.240)	(0.157)
Network HP growth, %	−0.007	0.077	0.120	−0.212
	(0.063)	(0.117)	(0.198)	(0.215)
Years	08-17	08-17	08-09	10-17
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Observations	49,703	13,072	2,356	10,372
R-squared	0.244	0.158	0.527	0.177
Note:			*p<0.1; **p<0.05; ***p<0.01	

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

tablishments themselves, rather than their location, were important determinants of employment.

On the other hand, the baseline weighting I use might capture the fact that nearby ELAs with a large number of employees are likely to serve more customers and therefore the demand shock in those areas should be more important. In either case, the magnitudes are not too dissimilar and the main thrust of the argument remains intact.

In the last set of tests I exploit the Irish Sea to investigate the importance of establishment networks in Northern Ireland. This test is slightly less clean than my Home County test because Northern Ireland is different from Great Britain in a number of ways. Its housing market dynamics are much closer to that of Ireland and a relatively small number of locally non-tradable firms have links to the mainland. That said, exploiting the fact that people cannot easily cross the

Table 3.10: Large Firms with Equal Weighting

	Locally non-tradable employment annual growth			
	(1)	(2)	(3)	(4)
	Within 50km	Beyond 50km	Within 100km	Beyond 100km
Local HP growth, %	−0.047	−0.053	−0.040	−0.055
	(0.049)	(0.047)	(0.048)	(0.047)
Network HP growth, %	0.240***	0.506	0.321***	0.072
	(0.062)	(0.307)	(0.071)	(0.183)
Years	10-17	10-17	10-17	10-17
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Observations	243,698	251,977	250,683	251,410
R-squared	0.144	0.143	0.143	0.142
Note:			*p<0.1; **p<0.05; ***p<0.01	

All regressions are weighted by employment and standard errors (in parentheses) are double clustered at the firm and local-authority level. Large firms are defined as having more than 1,000 employees and the network term is a simple average of house price changes in networked local authorities.

sea to visit restaurants or get their cars fixed should shed light on whether there might be non-spatial implications of establishment networks.

Table 3.11 shows results analogous to those in Table 3.6. Across the four columns I look at locally non-tradable firms that have a presence in Northern Ireland. In the first two columns we see that the network coefficient is statistically significant for large firms but not for those employing fewer than 1,000 people. In the second column, where I put a zero weight on ELAs in Great Britain, the network term increases in magnitude but becomes more imprecise. In the final column, where zero weights are placed on other Northern Irish ELAs, the network coefficient turns strongly negative. At face value, that suggests a negative demand shock in Wales might lead to an increase in employment in Northern Ireland. This result is a slight puzzle as the magnitude of the coefficient is relatively large. But it is likely explained by the disjointed housing markets and economic structures

Table 3.11: Establishment-Level Network Eects in Northern Ireland

	Locally non-tradable employment annual growth			
	(1)	(2)	(3)	(4)
	Small firms Full network	Large firms Full network	Large firms exc. GB network	Large firms only inc. GB network
Local HP growth, %	0.151	−0.134	−0.106	−0.036
	(0.133)	(0.077)	(0.085)	(0.089)
Network HP growth, %	−0.104	0.419***	1.123**	−2.072***
	(0.131)	(0.077)	(0.447)	(0.645)
Years	08-17	08-17	08-17	08-17
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Observations	9,564	6,627	6,389	6,388
R-squared	0.160	0.136	0.136	0.135
Note:				*p<0.1; **p<0.05; ***p<0.01

All regressions are weighted by employment and standard errors (in parentheses) are double clustered at the firm and local-authority level. Large firms are defined as having more than 1,000 employees. For the last two columns the network term is calculated by taking a weighted average of networked local authority house price changes as before and setting weights to zero if the local authority is beyond or within Great Britain, respectively.

between Northern Ireland and the mainland. In any case, it does not provide convincing evidence for the alternative hypotheses of collateral constraints or information spillovers. The overall message from Table 3.11 therefore supports my case for local demand spillovers being more important than non-spatial network effects.

3.4.3 Council Tax Demand Shocks

House price changes are one way of capturing local demand shocks but they are by no means the only source of variation for household finances. In this part I use variation in English local authority council tax rates over time to investigate the importance of local demand and establishment networks for locally non-tradable firms. Council tax is levied at the local-authority level and, for ease of comparison, I use data on the 325 English local authorities between 2007 and 2011. The independent variable is simply the annual growth rate in the average council tax bill for that local authority. More details can be found in Appendix D.

Unlike changes in house prices, the advantage of using council tax shocks is that almost all households are affected because it is paid by outright owners, mortgaged owners and renters, with only a few exemptions. Council tax is also regressive as it accounts for a higher proportion of income for poorer households. Together, these features mitigate the drawback that changes in council tax amount to relatively small amounts of money (especially compared to multiple percentage point changes in asset prices worth hundreds of thousands of pounds).

Council tax is often used as the marginal source of local government funding but public services are unlikely to be good substitutes for locally non-tradable consumption, so I treat it as an exogenous income shock for households. One concern with this approach is that local governments might levy taxes in response to, or anticipation of, local economic conditions. The exogeneity condition therefore relies on the assumption that most local government budgets are not closely related to the demand for locally non-tradable goods and services.[26]

Table 3.12 shows the council tax regressions, which are analogous to those in the main set of results.[27] In general, the results convey a remarkably similar overall

[26] This does not seem unreasonable. Other documented things that determine local council tax rates are the (often criticised) national distribution of resources using the Barnett Formula (e.g., McLean and McMillan (2003)); spatial mimicry between neighbouring local authorities (e.g., Revelli (2001)); and the interaction between local and national politics (e.g., James and John (2007)).

[27] Due to complications stemming from the exact timing of the council tax shocks, in the following specifications I use contemporaneous council tax shocks in the weighting of the network term.

Table 3.12: Council Tax Specification

	Locally non-tradable employment annual growth			
	(1)	(2)	(3)	(4)
	Small firms Full network	Large firms Full network	Large network Full network	Large network Network beyond 100km
Local council tax shock, %	−0.008	−0.212***	−0.324***	−0.297***
	(0.040)	(0.054)	(0.065)	(0.068)
Network council tax shock, %	−0.001	−1.718	−12.104***	−3.771***
	(0.070)	(1.525)	(1.956)	(1.315)
Years	08-11	08-11	08-11	08-11
Fixed effects	Yes	Yes	Yes	Yes
Specification	OLS	OLS	OLS	OLS
Observations	90,115	102,497	43,683	43,684
R-squared	0.307	0.261	0.283	0.279
Note:	*p<0.1; **p<0.05; ***p<0.01			

All regressions are weighted by employment and standard errors (in parentheses) are double clustered at the firm and local-authority level. Large firms are defined as having more than 1,000 employees and large-network firms are defined as having a network of more than 250 ELAs.

message about the importance of local demand shocks and their propagation. The first column shows there is little evidence that local or distant council tax shocks affect employment decisions of firms employing fewer than 1,000 people. Once again the main effect comes through larger firms and the second column suggests that a 10% increase in local council tax leads to around a 2pp fall in locally non-tradable employment growth. The network term is of plausible magnitude and direction but is statistically insignificant (the signs are reversed because positive tax shocks detract from demand). When defining firm size by the number of ELAs the firm is connected to in the third column, the local effect increases slightly to around 0.3 and the network term is a statistically significant -12. Again dividing by the 389 connections, that is roughly an tenth of the size of the local demand effect, which is a little smaller than the relative size of the network term using house price shocks. In the last column I put a zero weight on ELAs that are less than 100km away and I find the network effect falls by almost two thirds. Once again, this is consistent with the bulk of the network term being explained by modest spatial demand shock spillovers rather than information or balance sheet constraints.

3.5 Conclusion

The structure of firms has historically often been abstracted away when identifying firm shocks and forming narratives around propagation channels. I use establishment-level employment data to investigate the importance of spatial demand shocks and within-firm propagation for firms that generate revenues from people who reside nearby. I show that house prices changes are associated with economically significant employment semi-elasticities for these locally non-tradable firms. But I find no convincing evidence that these shocks propagate through balance sheets or other non-spatial linkages within firms' establishment networks. Instead, I show that the most convincing reason that demand shocks matter is that people travel across administrative borders; demand shocks ripple across the countryside.

In many ways this is a surprising result because it goes against economic intuition that collateral constraints might play a role in the propagation of shocks. It is also counter to the findings in Giroud and Mueller (2017b) for the US. But there are two reasons for the apparent discrepancy. First, collateral constraints are likely to affect small firms that find it difficult to access diverse sources of funding.

For these firms, the intensive margin of employment adjustment is unlikely to capture the full extent of shock propagation. My unbalanced panel is much more suited to understanding the dynamics of larger firms. The result that local (and to some extent regional) demand shocks seem to matter for these larger firms furthers our understanding about how this segment of the economy operates. There is work to be done to fully capture small-firm establishments that shut down, delay opening or change qualitative employment contracts in the face of changing economic conditions.

Second, collateral constraints are likely to bind more when firms are faced with negative demand shocks. The heart of this chapter focusses on the period between 2010 and 2017, during which house prices were sluggishly recovering from their lows across most regions. My results suggest that spatial spillovers are likely to be relatively important when there are small positive local demand shocks. That tells us something about how the economy operates in normal times. But I also show that the magnitude of employment responses may be larger in the face of negative demand shocks.

This work suggests that house price changes and other local demand shocks might have a slightly wider impact than otherwise expected. Using the establishment network of UK firms, it is possible to estimate where future shocks might do their most damage. For example, a large employer laying off its geographically-concentrated workforce is likely to affect the employment of its home local authority through direct and second-round, local-demand effects. Those local-demand spillovers might reach modestly beyond the immediate administrative boundaries.

Chapter 4

The Granular Hypothesis Revisited: Demand or Supply?

This chapter uses the universe of employment data to quantify the size distribution of firms and establishments in the UK. Although firm size can be modelled as a Pareto distribution with a quantitavely fat tail, I show this is primarily driven by the presence of a small number of large locally non-tradable businesses. These firms exhibit low geographical concentration and have expansive networks of establishments across the country. Given I find limited demand spillovers between these establishments in Chapter 3, I show that the propagation mechanism behind the Granular Hypothesis is more likely to work through supply-side disruptions.

4.1 Introduction

Many phenomena outside the natural world exhibit fat tails. In particular, the existence of a large number of small firms and few very large ones is readily apparent to the casual observer, yet interesting enough for researchers to spend time analysing and attempting to quantify. In this chapter I will exploit the universe of UK employment data to challenge the way we think about firm-size employment distributions and their importance for the propagation of shocks in the economy.

The Granular Hypothesis popularised by Gabaix (2011) showed, in the presence of fat tails, idiosyncratic shocks to the largest firms in the economy can account for a significant proportion of aggregate volatility. Although this has been challenged empirically by testing the relative contribution of firm-level shocks (e.g., Foerster, Sarte, and Watson (2011) and Stella (2015)), there has been little attempt to investigate the properties of large firms, or consider the channels through which the Granular Hypothesis might operate.

I document that the firm-size distribution in the UK does indeed have a quantitatively fat tail and this is largely driven by the presence of firms that generate revenues from customers who live nearby. In the previous chapter I found there is limited evidence for the propagation of local demand shocks across firm establishment networks. In this chapter I find the size distribution of these locally non-tradable firms is closely related to the number of establishments in their network (i.e., they exhibit low geographical concentration). Together, these facts imply that supply shocks are likely to be more important than demand shocks in explaining the Granular Hypothesis.

I first explore under what conditions it is appropriate to refer to the firm-size distribution in a quantitative sense. All fitted distributions are approximations to the true data. In the case of the UK, measurement error aside, the true data are known with certainty. This has implications for how the data should be described. In some cases, quantifying approximate distributions can usefully convey information about the relative importance of large firms. But this is conceptually different from fitting a probability distribution to a series of unknown states of the world.

Second, I probe whether we should follow a long line of researchers (e.g., Axtell (2001) in the case of firm size) and fit a Pareto distribution to the data. The fact that Pareto distributions have appealing theoretical properties is a separate issue from understanding what a calibrated distribution might imply about real world phenomena, such as implications of the Granular Hypothesis. I show that Pareto distributions are useful to help us write down models and quantify the difference between different subsamples of data, but they should rarely be thought of as the true (or even the best-approximating) distribution. In many cases, lognormal distributions are almost certainly more likely to yield better fits in a statistical sense but might make it harder to convey information.

Third, I show that the exact parametrisations of the fitted Pareto distribution can be sensitive to a whole range of issues. Broadly speaking, however, the UK firm-size distribution exhibits a quantitatively fat tail in the sense that a reasonably calibrated Pareto distribution has an undefined variance. Within that, locally non-tradable firms have a high concentration of very large employers. Although they only account for around a fifth of private-sector employment, more than half of the 50 largest employers are locally non-tradable firms. The fat tail of UK firms is therefore mechanically driven by the firm-size distribution of locally non-tradable firms.

Finally, I consider subsets of the universe of UK employment and the implications for the presence of fat tails for the propagation of macroeconomic shocks. I show that the firm-size distribution of locally non-tradable firms is fat because constituent firms have a large establishment network, in contrast to locally tradable firms that concentrate employment in a few locations (e.g. large manufacturing plants). That means the largest firms in the UK by employment are spread all over the country. Given the findings in the previous chapter that local demand shocks do not propagate meaningfully across firm establishment networks, and that economy-wide preferences only shift slowly over time, that suggests supply chain linkages are likely to be the key mechanism behind the Granular Hypothesis.

For a long time there was insufficient data to be confident about the exact skew of the UK employment distribution. Notwithstanding, the East India Company at the turn of the nineteenth century was almost certainly on the same relative scale as modern day multinationals, and British railway companies employed almost 5% of the labour force by 1850.[1] It is likely that firm-size distributions have looked qualitatively similar for generations but for a long time it was assumed not to

[1]See Richards (2002) and Bowen (2005), respectively.

matter. The view that idiosyncratic shocks were first-order irrelevant because of the law of large numbers, as set out in Lucas (1977), was one of the main analytical motivations for developing representative agent models in the 1980s. A natural corollary was that distributions of agents, shocks and economic linkages had little relevance for the macroeconomy.

That view is now recognised as being incomplete. From as far back as Long and Plosser (1983), Horvath (1998) and Horvath (2000), serious contributions have been made to show how the law of large number breaks down in the presence of fat-tailed distributions and this was brought into the modern day debate with the seminal work of Gabaix (2009). Fat-tailed distributions ensure microeconomic shocks might matter for aggregate volatility, perhaps because of the structure of supply chains (e.g., Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012); Acemoglu, Akcigit, and Kerr (2016); Baqaee (2016); and Baqaee and Farhi (2017)); or because the process that governs macroeconomic tail risks itself has an elongated tail (e.g., Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017)). This chapter fits into this literature by reconsidering whether the firm-size distribution does indeed have a fat tail. I argue that, to the extent the Granular Hypothesis explains much aggregate volatility at all, we should focus our attention on subsets of firms and their economic structures, which is consistent with Grassi (2018).

Section 2 summarises some of the key UK employment data and Section 3 introduces the Pareto distribution, which is common when working with data of these sort. Section 4 is the heart of the chapter and explores the four questions I outline above. Section 5 concludes.

4.2 Top-Down Statistics

Using BSD data, in 2017 there were about 2.6 million non-public enterprises operating in the UK, running 2.9 million establishments and employing around 24 million people.[2] Around 43% of firms only employ one individual, which highlights the high concentration of small firms that has long been recognised as an empirical regularity but was emphasised more recently Gabaix (2009). On average, UK firms employ 9 people but this is heavily skewed by the numbers of single-employee firms. Conditional on having at least two employees, UK firms have an average of around 15 workers. At the other end of the size distribution, conditional on having at least 1,000 employees, firms have an average of 4,600 workers. The largest 1,000 firms employ a total of 7 million people and the largest 10,000 firms employ 11 million. This is summarised in Table 4.1.

Table 4.1: Private-Sector Employment by Firm Size

Employment	No. enterprises	Cumulative enterprises	Total employment	Cumulative employment
1	1,129,394	43.4%	1,129,394	4.7%
2 to 5	1,009,072	82.2%	2,792,678	16.4%
6 to 10	224,171	90.8%	1,681,937	23.5%
11 to 20	122,260	95.5%	1,786,385	30.9%
21 to 50	73,954	98.3%	2,312,983	40.6%
51 to 100	23,485	99.2%	1,638,200	47.5%
101 to 250	11,986	99.7%	1,842,164	55.2%
251 to 500	3,897	99.9%	1,356,477	60.9%
501 to 1,000	1,904	99.9%	1,322,570	66.4%
1,001 to 5,000	1,417	100.0%	2,913,664	78.6%
More than 5,000	321	100.0%	5,114,212	100.0%
Total	2,601,861		23,890,664	

This table covers all private-sector employment in 2017 from the BSD.

[2]There were around 750,000 inactive and 250,000 non-private establishments in the data that have been excluded.

Much of the analysis in this thesis splits UK firms into those that rely on nearby customers for their revenues and those that do not. Understanding the size distributions of these subsets is intuitively interesting because the former often exhibit a low geographical concentration. In other words, firms that serve their local community (defined as locally non-tradable firms in Chapter 1) tend to have a large number of establishments relative to the total employment of the firm. Using the mapping described in Appendix A, just over a fifth of private-sector employees currently work in the locally non-tradable sector. In 2017 there were around 360,000 locally non-tradable enterprises, operating 500,000 establishments and employing 5.4 million people, as shown in Table 4.2. Locally non-tradable firms tend to be larger than other types of firm, with an overall average size of 17 people. Conditional on having at least two employees, locally non-tradable firms employ an average of 21 people. More importantly, locally non-tradable firms are well represented among the UK's largest firms. Ten enterprises in the top twelve largest employers in the UK can be categorised as locally non-tradable, and 28 are in the top 50 largest employers.

The sample I use in the previous chapter runs from 2008 to 2017 and the broad characteristics of UK employment were relatively constant over that period. The locally non-tradable employment share has stayed relatively constant over time, though in 2010 there were 4.7 million people employed in the locally non-tradable sector out of a private-sector workforce of 23 million. Between 2008 and 2017 the total number of enterprises and establishments increased by about 360,000

Table 4.2: Locally Non-Tradable Employment by Firm Size

Employment	No. enterprises	Cumulative enterprises	Total employment	Cumulative employment
1	83,996	23.1%	83,996	1.6%
2 to 5	175,824	71.4%	541,950	11.6%
6 to 10	57,437	87.2%	429,860	19.6%
11 to 20	27,915	94.9%	404,000	27.1%
21 to 50	13,461	98.6%	401,974	34.6%
51 to 100	2,797	99.4%	193,819	38.2%
101 to 250	1,257	99.7%	191,195	41.8%
251 to 500	489	99.8%	171,618	45.0%
501 to 1,000	251	99.9%	177,038	48.2%
1,001 to 5,000	244	100.0%	507,925	57.7%
More than 5,000	95	100.0%	2,276,057	100.0%
Total	363,766		5,379,432	

This table covers locally non-tradable employment in 2017 from the BSD using the definitions in Chapter 1 and Appendix A.

and was mainly driven by an increase in single-person firms.[3] These small net changes hide the substantial turnover of registered firms during this period but it is common for advanced economies to have significant rates of firm entry and exit. In many cases, although the administrative records show a firm death followed by a birth soon after, its economic function in terms of employment and output remain relatively similar.[4]

4.3 The Firm-Size Distribution

Across the world, employment is usually far from uniformly distributed across businesses. Beyond talking loosely about large numbers of small firms it is often helpful to characterise the shape of a distribution by quantifying parameters or describing statistical moments. A small number of large firms means the firm-

[3] The number of single-person firms increased from 860,000 in 2008 to 1.13m in 2017. Much of that increase was likely driven by tax incentives.

[4] This is especially true for the restaurant and entertainment sector. Whether or not this matters depends on the question at hand. If the same people are fired and then rehired with a change in ownership that might tell us something about short-term frictional unemployment but less about medium-term employment trends.

size distribution is heavily skewed to the right: the modal firm is smaller than the median firm, which is in turn smaller than the mean.[5] One mathematical function that shares some of these empirical regularities is the Pareto distribution. The exact functional form of the Pareto probability distribution, where s is firm size and s_0 is the minimum size cut-off, is defined as

$$p(s) = \frac{\alpha}{s_0}\left(\frac{s}{s_0}\right)^{-(\alpha+1)} \tag{4.1}$$

When applying this to firm size distributions, a discrete approximation is often used to characterise the upper tail. In this case, the counter-cumulative distribution function is

$$\Pr(S > s) = \int_x^\infty p(s)ds = \left(\frac{s}{s_0}\right)^{-\alpha},\ s \geq s_0 \tag{4.2}$$

In Equations 4.1 and 4.2, the parameter α captures how extended (or fat) the tail is. In the special case that $\alpha = 1$, Equation 4.2 is called a Zipf distribution and this is an oft-used characterisation of many phenomena in the natural and man-made world.[6] The most interesting feature of the Pareto distribution is that the mean only exists when $\alpha > 1$ and its variance is defined only in cases where $\alpha > 2$. Small values of α are therefore associated with a small number of very large observations that dominate the sample.

There are a number of ways of estimating the tail parameter. One common

[5]This study focuses on firm size measured by the number of employees but many of distributional features are similar whether firm size is measured in number of employees or turnover.

[6]For example, see Gabaix and Landier (2008); Luttmer (2007); Fujiwara (2004)

approach is to take logs of both sides of the discretised version of Equation 4.2 and estimate α using ordinary least squares (OLS).[7]

$$\log p(s) = \log C - (\alpha + 1) \log s + \varepsilon \tag{4.3}$$

An alternative approach first devised by Hill et al. (1975) is to take the maximum likelihood estimate

$$\mathcal{L} = \ln \prod_{i=1}^{n} \frac{\alpha}{s_0} \left(\frac{s_i}{s_0} \right)^{-(\alpha+1)} \tag{4.4}$$

and set $\frac{\partial L}{\partial \alpha} = 0$ to get

$$\hat{\alpha} = n \left[\sum_{i=1}^{n} \ln \frac{s_i}{s_0} \right]^{-1} \tag{4.5}$$

In perhaps the most famous US study documenting the empirical size distribution of companies, Axtell (2001) uses an OLS regression for firm size buckets and estimates a tail index, $\hat{\alpha}$, of 1.06. If firm size is normally distributed then it is well known that aggregate output fluctuations should be proportional to $1/\sqrt{N}$, where N is the total number of firms (e.g., see Gabaix (2011) and Carvalho and Grassi (2015)). In fact, this is true as long as $\alpha \geqslant 2$. If firm size is distributed with a Zipf law, however, volatility decreases at the significantly slower rate of $1/\log N$. For the intermediate case where $1 < \alpha < 2$, volatility is proportional to $1/N^{1-\frac{1}{\alpha}}$. The fact that Axtell's estimated tail coefficient is below 2 suggests that the firm-size

[7] This estimation is often done in a discrete fashion by ranking firm size S from largest to smallest and running a rank-size regression. For example, see Axtell (2001).

distribution exhibits a fat tail, in the sense that the law of large numbers does not strictly apply, and so idiosyncratic shocks to the largest firms are more likely to lead to aggregate consequences. To put this in context using data from Table 4.1, if α is estimated above 2, volatility is proportional to 1/1613 in the standard models. If the estimate of α is below 1 then volatility is proportional to 1/14.8, which is two orders of magnitude larger.

4.4 Fitting Pareto Distributions to the Data: Four Questions

In this section my aim is to challenge the casual use of the Pareto distribution and explore whether it is appropriate to classify the UK firm-size distribution as having a fat tail in a quantitative sense. This matters because firm-idiosyncratic shocks can have macroeconomic consequences in Gabaix (2011). Understanding the properties of the firm-size distribution therefore helps us understand the shocks and transmission mechanism that might be pertinent for aggregate output or employment volatility. For the remaining calibrations I use 2017 data but the results are similar across time.

4.4.1 Why Fit Mathematical Distributions?

Care must be taken when attempting to fit mathematical distributions with an infinite domain to the universe of firms, even if it is a useful approximation to help answer certain questions. Using tractable functional forms can be helpful when developing models that come close to matching the distribution of firm size but also shed light on another part of the economy (e.g. Melitz and Ottaviano (2008)). But it needs more consideration when the exact functional form is important for the result in question, such as in Gabaix (2011).

One field that often uses continuous distributions to fit data is finance, where market movements are inherently uncertain and experience has shown that asset price movements are far from normally distributed. The use of tractable functions can help us understand the unknown probability distribution associated with various future states of the world, which can then be used to inform policy (e.g., Cumming and Noss (2013)). Importantly in these studies, the domain of the true probability distribution is infinite. No matter how much data we can collect, there will always be a role for approximating a probability distribution over future events, even if that distribution collapses to a point mass.

But things are slightly different when summarising cross-sectional phenomena. Gabaix (2009) rightly points out that using a Gaussian distribution to approximate the distribution of population height is often a sensible thing to do even though it only applies to part of the domain (since heights cannot be negative). But such calibrations tend to be robust to small perturbations in parameters and have little consequence for our understanding for broader questions in science or society. Gabaix (2011) goes on to suggest that fitting a distribution to UK firm-size data might be appropriate for an economist who knows the UK's GDP and wants to estimate roughly what kind of firms might exist.[8] That is a somewhat

[8]See p. 741 of Gabaix (2011)

redundant exercise to an economist with access to the universe of UK data. All population moments exist, are known with precision and the probability of finding a firm larger than the largest in the sample is precisely zero. More generally, probability and cross-sectional-summary distributions are conceptually different and should be used for different purposes.

Fitting functional forms to the latter can be helpful because they convey information about the relative importance of different sized firms. That is challenging to do using qualitative descriptions alone. But we need consider what a quantitative calibration can, and cannot, tell us. Showing that one country's estimated tail parameter is larger than another might indicate the relative importance of small shocks to the largest firms. But it seems less appropriate to mechanically apply probability principles or take parametrisations too literally given the conceptual problems involved.

4.4.2 Why Calibrate a Pareto Distribution?

Let us continue under the assumption that fitting mathematical distributions to the data is a useful exercise. Which family of distributions should we pick? Pareto distributions have been used to characterise economic phenomena for many years because they are easy to work with. The simple functional form and the useful mathematical properties have long made it irresistible for those wanting to develop tractable models to investigate how firm production might impinge upon the macroeconomy (e.g., Houthakker (1955)). Slightly later, the seminal example of Lucas (1978) showed how the size distribution of firms is a simple transformation of the managerial talent distribution when the latter is modelled as a Pareto distribution. More recently, Geerolf (2016) showed that Pareto distributions can arise from micro-founded production functions. The international trade literature has increasingly turned to the Pareto distribution to introduce agent heterogeneity and overturn established results (e.g., Melitz and Ottaviano (2008); Chaney (2008); and Eaton, Kortum, and Kramarz (2011)).[9]

On the empirical front, a caricature of many studies that attempt to fit distributions to various phenomena goes as follows. First the researchers obtain a data set that appears to exhibit a fat tail. Next they try to fit a lognormal distribution but decide to add one or more Pareto distributions to subsets of the data to improve the fit (e.g., see Reed and Jorgensen (2004); Fazio and Modica (2012); and Luckstead and Devadoss (2017)). Often the motivation for such an approach is that simple OLS regressions only hold for parts of the data or that we learn more by approximating the true distribution with more and more complex piecewise or convoluted functions. Such research might not be without merit in general but I argue is unlikely to be especially helpful when drawing conclusions for issues such as the propagation of shocks in the economy.

It is certainly true that one can find a calibrated version of the Pareto distribution that is not visually inconsistent with the US or the UK firm-size distribution. As Axtell (2001) shows, by grouping firms into buckets and plotting the log frequency against log firm size, the data points line up in a relatively straight line, which is required for the Pareto distribution fit. But the Pareto distribution is

[9] But more recent work by Bas, Mayer, and Thoenig (2017) shows that Pareto distributions might lead to incomplete characterisations of trade relationships.

just one class of many functional forms that can fit a data set with a small number of very large observations and a large number of small ones. Axtell does not test the fit of calibrated distributions from other families (visually or econometrically), let alone whether the Pareto distribution is the best fit in some objective sense.

An alternative distribution that can be used to model firm-size data is the lognormal distribution, which has a finite mean and variance for all parametrisations.[10] Figure 4.1 shows a stylised example of the two classes of distribution. On the log-log plot, the Pareto distribution is a straight line and its slope determines which moments are finite. The shallower it is, the fatter its tails. The lognormal distribution is always concave to the Pareto distribution, so might look like a straight line from a distance but it will always have a increasingly negative second derivative (and this ensures its variance is finite).

Figure 4.1: Stylised Distributions

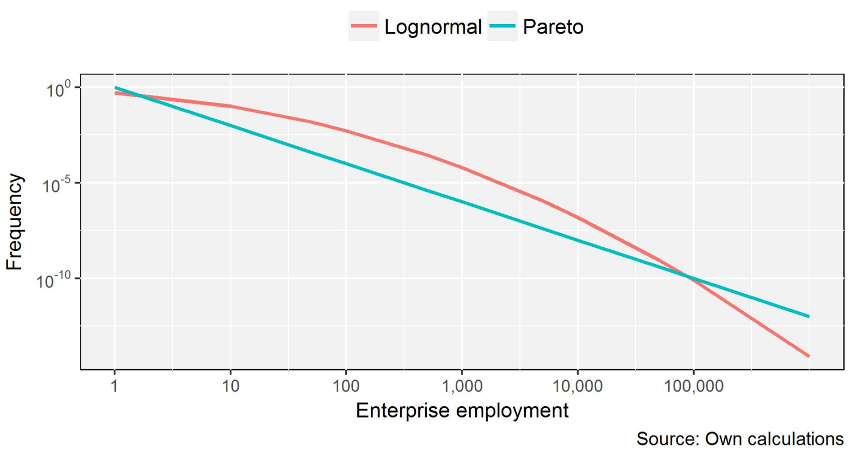

This figure shows two example distributions. The lognormal distribution is is calibrated with $\mu = 0$ and $\sigma = 1.8$. The Pareto distribution is calibrated with $s_0 = 1$, $c = 1$ and $\alpha = 1$.

[10]For a mean of μ and a variance of σ^2 of the variable's natural logarithm, the mean of the distribution is $\exp(\mu + \frac{\sigma^2}{2})$ and the variance is $(\exp(\sigma^2) - 1)\exp(2\mu + \sigma^2)$.

Axtell notes there are too few very small and very large firms in the US data compared to the Pareto fit, and the same is true in the UK. This is unsurprising because of the discrete nature of the data but the implication is that a lognormal distribution is likely to fit the data at least as well as a Pareto distribution (since too few large observations imply they lie below the straight line on a log-log plot).

One way to test which distribution fits the data better is to use a likelihood ratio test based on the Kullback-Leibler criteria of relative entropy.[11] Figure 4.2 below shows the p-values for this test at different minimum size cut-offs. The null hypothesis is that both the Pareto and the lognormal distribution fit the data equally well. The alternative hypothesis is that the lognormal distribution fits the data better based on the argument above. The black line shows that we can reject the null hypothesis at the 5-percent level for most values of the firm size cut-off below around 2,000 employees (i.e. for most reasonable values of s_0). This suggests that, when faced with a binary choice, a lognormal distribution might fit UK employment data better than a Pareto distribution.

We know from Table 4.1 that firms employing fewer than 1,000 people account for the vast majority of firms but only around two thirds of total employment. The extended x-axis in Figure 4.2 shows that there is less evidence against a Pareto fit for larger values of the firm-size cut-off. This highlights the importance of choosing an appropriate value for s_0.

[11] See methodology proposed by Vuong (1989) for more details. The test statistics is the sample average of the standard-deviation standardised average of the log-likelihood ratio of the two distributions and follows a Normal distribution.

Figure 4.2: Lognormal and Pareto Distribution Likelihood Ratio Test

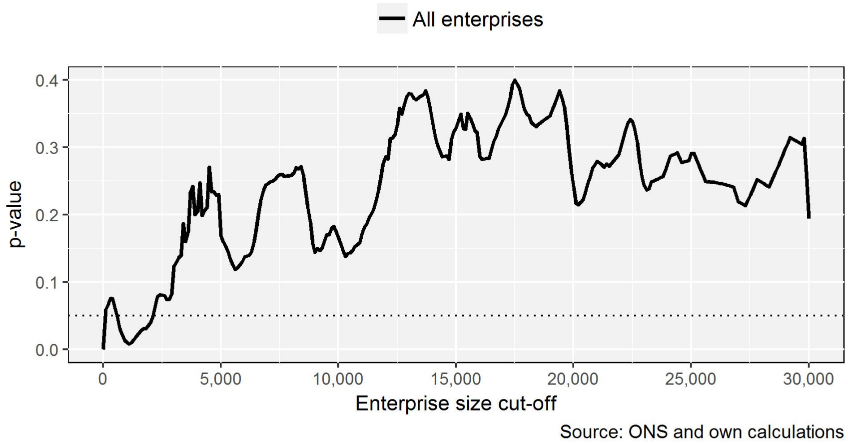

This figure shows the p-values for likelihood ratio test based on relative entropy. The null hypothesis is that both distributions fit the data equally well. The alternative hypothesis is that a lognormal distribution fits the data better.

The evidence presented in Figure 4.2 suggests we should be slightly careful when fitting Pareto distributions to firm-size data, even if it is the most common starting point for analysis of this kind. But the advantage of using the Pareto distribution is that it can still give us information about α and therefore a sense of how fat-tailed the data are.[12] A low α means the very largest firms might have a significant role to play in generating aggregate fluctuations. We should nevertheless be cautious about declaring that the Pareto distribution unequivocally provides the best fit for the upper tail of firms.

[12]For example, if there a a small number of very large firms, a lognormal distribution might fit the data points in the flat part of the curve before begins to slope down. In this case using a Pareto distribution and estimating the appropriate α for a larger value of s_0 would still tell us something about the relative importance of those largest firms.

4.4.3 Which Pareto Distribution?

Axtell's estimated α of 1.06 suggests that the US firm-size distribution exhibits a quantitatively fat tail in the sense that the variance of the approximating distribution is undefined. Using the BSD, I find almost identical estimates for α in 2017 (see the first row of Table 4.3) and the overall data show a similar pattern. This is a useful benchmark but there are well-known problems with using simple regression coefficients to estimate the tail parameter (see Clauset, Shalizi, and Newman). First, regression analysis requires the minimum size cut-off to be set in an ad hoc way even though results can be quite sensitive to it. Second, it is difficult to get a sense of how well the estimated α fits the data, even with the appropriate cut-off. Standard errors are incorrect because the logarithmic structure breaches the Gaussian noise assumption and the r^2 can often be very close to 1 even if the fit for the most important (large) firms is poor.

For the rest of the distributional analysis I therefore put more weight on maximum likelihood estimates as this should provide slightly better estimates than the regression outputs. Figure 4.3 shows the Hill estimator for $\hat{\alpha}$ at different firm size cut-offs. The black line shows that when the universe of firms is included in the analysis the estimated tail parameter starts off close to 1 (consistent with the regression coefficients and Axtell (2001)) and remains well below 2 across all cut-offs. When including all firms, the firm size distribution appears to have a fat tail.[13]

[13]An alternative measure of large-firm dominance (and hence firm concentration) is the Herfindahl index. The 50 largest locally-non tradable and locally tradable firms in 2017 had a Herfindahl index of around 0.05 and 0.03, respectively. Grassi (2018) shows this summary statistic is directly related to the propagation of firm-level shocks in his analytical model. But

Figure 4.3: Estimated Tail Parameters

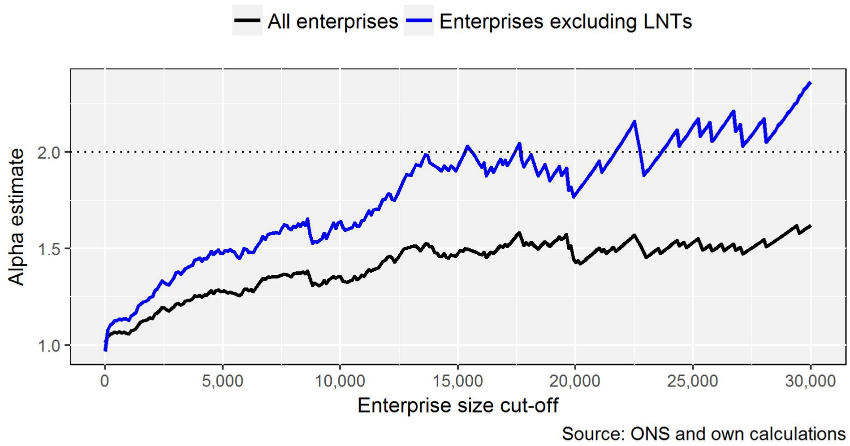

This figure shows the maximum likelihood estimate of the tail parameter for all firms (black) and firms excluding locally non-tradable (LNT) businesses (blue) at different size cut-offs.

The blue line shows the results when locally non-tradable firms are dropped from the analysis. Once again, the estimated tail parameter starts close to 1 but this time rises slightly quicker. When the size cut-off reaches firms with more than around 15,000 employees, it approaches 2, which implies the distribution has a finite variance and a thin tail above this point. This suggests that, to the extent UK firms as a whole might have a fat tail, it might in fact be driven by the presence of a few very large locally non-tradable firms. Of course, Table 4.2 shows there are fewer than 100 firms that employ more than 5,000 people so the estimates in Figure 4.3 are likely to be both imprecise and very specific on the exact interval considered.

Another way to visualise the importance of firm size and the difference between the employment structures across sectors is to plot the cumulative employment

just as there are problems defining the firm-size cut-off when calibrating firm-size distributions, so too are there problems regarding which sectors, or how many firms, to include in a Herfindahl index calculation.

accounted for by firms employing different numbers of employees. Figure 4.4 makes clear that only half of locally non-tradable employment is accounted for by firms employing fewer than 1,000 people. In stark contrast, firms smaller than 1,000 people make up almost 80% of locally tradable employment. The fact that the blue line is so far below the red line in Figure 4.4 shows that locally non-tradable firms are relatively dominated by a small number of large firms.

Figure 4.4: Total Employment by Firm Size

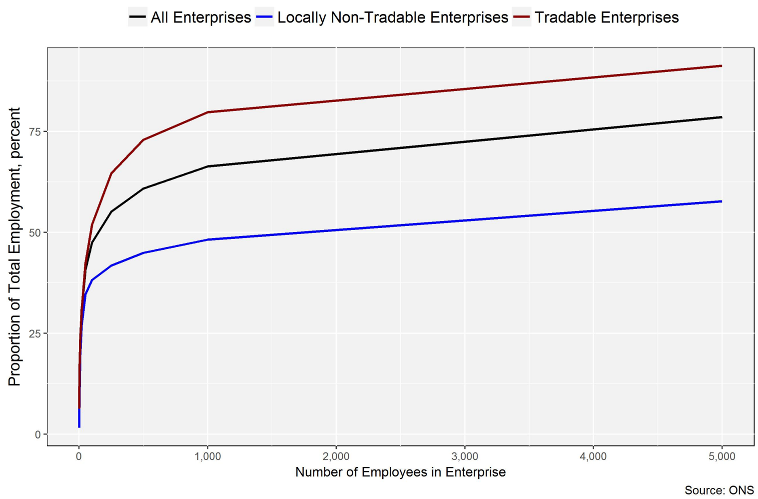

This Figure shows the cumulative employment by firm size for different subsets of enterprises. Firms larger than 5,000 employees are excluded to preserve anonymity. Total employment refers to private-sector employment only.

Returning to quantifying firm-size distributions, there are essentially three ways to estimate the tail parameter, α. One can choose a value for s_0 and then either run OLS or MLE on the subsample of firms larger than the size cut-off. In Table 4.3 the first few rows show these estimates when I take the largest 0.1% of observations. One drawback of these two approaches is that is difficult to point to a sound theoretical justification for any ad hoc prior for the cut-off. The third method uses MLE to simultaneously estimate the size cut-off and the tail parameter. This is done in the last few rows and I allow it optimise over the largest 1% of observations. Despite aforementioned concerns with using OLS, all three methods yield broadly similar results.

Table 4.3: Parameter Estimates of Firm-Size Distributions

	Enterprises		
	All	Non-tradable	Tradable
Imposed cut-off			
OLS tail parameter, $\hat{\alpha}$	1.12	0.88	1.35
MLE tail parameter, $\hat{\alpha}$	1.06	0.82	1.33
Cut-off, s_0	700	900	650
Observations	2,600	360	250
Endogenous cut-off			
MLE tail parameter, $\hat{\alpha}$	1.05	0.83	1.29
MLE size-cutoff, $\hat{s_0}$	125	126	219
Observations	15,600	1,900	1,000

The first four rows provide estimates for the top 0.1% of observations. The last few rows jointly estimate the cut-off and tail parameter from the top 1% of observations. All reported observations and s_0 are rounded to preserve anonymity.

The first column confirms the presence of a fat tail in the UK firm-size distribution in that $\hat{\alpha}$ is close to 1 for all estimation techniques. These estimates are directly comparable to Axtell (2001).[14] The second and third columns perform the same analysis on locally non-tradable and tradable firms, respectively. Consistent with the message in Figures 4.3 and 4.4, locally non-tradable firms appear to be more dominated by a small number of very large firms and have a relatively fatter tail than tradable firms. But both estimates of $\hat{\alpha}$ are less than 2, so both distributions have a fat tail in an absolute sense.

The high-level answer to this question then is that the most plausible calibration of a Pareto distribution to fit UK data has an α somewhere between 1

[14] The results in Axtell (2001) implicitly assume a cut-off of 1; the cut-off only tends to matter deep into the tail.

and 2, which implies the law of large numbers does not hold because the variance of that distribution is undefined. The estimate of α increases the further into the tail one goes (i.e. the tail becomes thinner) but it is difficult to know exactly which calibration makes most sense to understand shock transmission. Within the overall distribution, firms reliant on local demand have a relatively fatter tail and are therefore likely to propagate shocks to the macroeconomy more than tradable firms.

4.4.4 Large or Many Establishments?

It is possible to run these tail-parameter estimates on any subset of the enterprise population. I have chosen to look at locally non-tradable firms because it gives us information about the potential mechanism at play in the Granular Hypothesis. For example, if the firm-size distribution matters for macroeconomic volatility, then there must be a reason why one large firm is fundamentally different from many small firms if they are performing the same economic function. Mathematically, the law of large numbers breaks down because idiosyncratic shocks do not wash out across firms. Economically, however, that must mean that constituent firms of a newly-formed conglomerate find that their shocks correlate after a takeover. Plausible explanations for that include the agglomeration leading to the concentration of managerial, supply-chain or financial risk.

One key difference between locally non-tradable and tradable firms is the structure of their establishment networks and this might give us an insight into why the firm-size distributions look different. Locally non-tradable firms require their customer base to be within close proximity to their establishments. To the extent that people are spread over the entire country, it seems likely that large locally non-tradable firms will be too.

Table 4.4 redoes the analysis in Table 4.3. This time the first three columns calibrate a Pareto distribution for the firm-size distribution when size is measured by the number of establishments in the firm's network. In the first column we see that there is now more uncertainty in the value of $\hat{\alpha}$, which is partly due to the fact that the endogenous cut-off estimation only wants to fit the tail to a few very large firms. Nevertheless, the message from the previous section is upheld with this new size definition: the law of large number fails to hold since α is likely below 2. The second and third columns also show a consistent message for locally non-tradable and tradable firms. Locally non-tradable firms are dominated by a small number of firms that have a large number of establishments, and their firm-size distribution has a fatter tail than the counterpart size distribution for tradable firms.

The final three columns of Table 4.4 ignore enterprise ownership and treat all establishments as independent units. This is tantamount to assuming each establishment takes little direction from head office and we can ignore the economic and legal bindings between establishments in the same firm. The fourth column shows the establishment-size distribution has an $\hat{\alpha}$ much closer to 2. More interestingly, the relative estimates of α are now reversed for locally non-tradable and tradable firms. This means that the largest locally non-tradable establishments are relatively smaller than those of tradable firms when compared to the rest of the relevant distribution.

Table 4.4: Parameter Estimates of Establishment-Size Distributions

	No. establishments			**Establishment employment**		
	All	Non-tradable	Tradable	All	Non-tradable	Tradable
Imposed cut-off						
OLS tail parameter, $\hat{\alpha}$	1.35	0.99	1.45	1.82	2.46	1.88
MLE tail parameter, $\hat{\alpha}$	1.12	0.83	1.60	1.77	2.94	1.33
Cut-off, s_0	170	50	5	500	400	670
Observations	2,600	360	250	2,900	500	260
Endogenous cut-off						
MLE tail parameter, $\hat{\alpha}$	0.98	0.84	1.66	1.82	1.86	1.63
MLE size-cutoff, $\hat{s_0}$	11	8	4	568	350	236
Observations	26,000	3,600	2,500	3,200	800	900

The first four rows provide estimates for the top 0.1% of observations. The last few rows jointly estimate the cut-off and tail parameter from the top 1% of observations. All reported observations and s_0 are rounded to preserve anonymity.

One way to visualise this reversal is to redo Figure 4.4 for the (counterfactually autonomous) establishments. Figure 4.5 shows that just over three quarters of people employed in the locally non-tradable sector work in establishments employing fewer than 125 people and almost everyone else works in establishments employing fewer than 1,000. In contrast, tradable establishments employing 125 account for just over half of tradable employment and around 20% of people work in establishments employing more than 1,000 people. In part, this is to be expected since the largest supermarket outlet is likely to be similar in size to the next few largest.[15] Customer catchment areas are naturally limited by willingness to travel and competition between firms. But the largest mines or factories are likely to be enormous given the economies of scale involved.

[15]There is limited publicly available information on the employment distribution across individual firms. But, the largest Tesco store in the UK has around 17,000 square metres of floor space compared to the average Tesco Extra, which has around 6,500 square metres. See Tesco PLC (2017) for more details.

Figure 4.5: Total Employment by Establishment Size

This Figure shows the cumulative employment by firm size for different subsets of enterprises. Establishments larger than 1,000 employees are excluded to preserve anonymity. Total employment refers to private-sector employment only.

Taken as a whole, Table 4.4 and Figure 4.5 therefore imply the firm-size distribution of locally non-tradable firms exhibits a fat tail because a small number of enterprises have a large establishment network across the country, even though each establishment is actually relatively modest in size. The firm-size distribution of tradable firms has a slightly less fat tail; there are a small number of large firms that concentrate their employment in very large establishments.

4.5 Conclusion

This chapter digs a little deeper into the UK firm-size distribution and extends the analysis of Gabaix (2009) and Gabaix (2011). The main section explores four questions related to estimating fat tails. First, I argue that it can be a useful exercise to fit continuous distributions to cross-sectional phenomena as long as the results are taken in context. In the above analysis, estimates of α tell us about the relative importance of large firms across subsamples of the data. This exercise is conceptually different from fitting a probability distribution to a series of unknown states of the world.

Second, I show that Pareto distributions are useful to help us write down models. In many cases, lognormal distributions fit the data better in a statistical sense but might make it harder to convey information about the data. Fitted Pareto distributions can be sensitive to a whole range of issues but, on the whole, the UK firm-size distribution exhibits a fat tail. This is driven by the firm-size distribution of locally non-tradable firms. In the final part I showed that the firm-size distribution of locally non-tradable firms is fat because constituent firms have a large establishment network, in contrast to locally tradable firms.

Two implications follow from these results. The first implication is that, when thinking about the failure of the law of large numbers and the propagation of small shocks through to macroeconomic volatility, we should focus on locally non-tradable firms. Although my estimates suggest the firm-size distribution of tradable firms also exhibits fat tails, locally non-tradable firms are likely to be more important in a mechanical sense.

The second implication is that, for macroeconomic propagation to be important, locally non-tradable establishments must surely be connected by more than mere association to an umbrella enterprise. If that is not true then the UK's fat tail can be broken down into its constituent establishments and α returns to a range where the law of large numbers holds.

Taken together, this gives us some information about which shocks might be especially susceptible for propagation. Work by Giroud and Mueller (2017b) suggests that local demand shocks might propagate across firm establishment networks but spillovers appear quantitatively small. In Chapter 3 I show this is especially true since the Great Recession in the UK. But locally non-tradable firms are often dependent on extensive supply chains and therefore work by Acemoglu, Akcigit, and Kerr (2016) and Baqaee (2016) is likely to be especially helpful in considering where vulnerabilities might develop. This chapter therefore provides motivation for further work to investigate the role of supply chain networks in locally non-tradable firms.

References

Acemoglu, D., U. Akcigit, and W. Kerr (2016): "Networks and the Macroeconomy: An Empirical Exploration," *NBER Macroeconomics Annual*, 30(1), 273–335.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80(5), 1977–2016.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017): "Microeconomic Origins of Macroeconomic Tail Risks," *American Economic Review*, 107(1), 54–108.

Adam, S., and J. Browne (2006): "A Survey of the UK Tax System," Discussion paper, Institute of Fiscal Studies.

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.

Arellano, M., and S. Bond (1991): "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *The Review of Economic Studies*, 58(2), 277–297.

Arellano, M., and O. Bover (1995): "Another Look at the Instrumental-Variable Estimation of Error-Components Models," *Journal of Econometrics*, 68(1), 29–51.

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.

Axtell, R. L. (2001): "Zipf Distribution of U.S. Firm Sizes," *Science*, 293(5536), 1818–1820.

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.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2016): "Social Networks and Housing Markets," Discussion paper, National Bureau of Economic Research.

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

Baqaee, D. (2016): "Cascading Failures in Production Networks," 2016 Meeting Papers 402, Society for Economic Dynamics.

Baqaee, D. R., and E. Farhi (2017): "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," Discussion paper, National Bureau of Economic Research.

Bas, M., T. Mayer, and M. Thoenig (2017): "From Micro to Macro: Demand, Supply, and Heterogeneity in the Trade Elasticity," *Journal of International Economics*, 108, 1–19.

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.

Bigio, S., and J. La'O (2016): "Financial Frictions in Production Networks," Discussion paper, National Bureau of Economic Research.

Blanchard, O. J., and L. F. Katz (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, 23(1), 1–76.

Bowen, H. V. (2005): *The Business of Empire: The East India Company and Imperial Britain, 1756–1833*. Cambridge University Press.

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.

Bunn, P., J. L. Roux, K. Reinold, and P. Surico (2018): "The Consumption Response to Positive and Negative Income Shocks," *Journal of Monetary Economics*, 96, 1 – 15.

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.

Carvalho, V. M., and B. Grassi (2015): "Large Firm Dynamics and the Business Cycle," .

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.

Chaney, T. (2008): "Distorted Gravity: the Intensive and Extensive Margins of International Trade," *American Economic Review*, 98(4), 1707–21.

Christelis, D., D. Georgarakos, T. Jappelli, L. Pistaferri, and M. Van Rooij (2017): "Asymmetric Consumption Effects of Transitory Income Shocks," .

Clauset, A., C. R. Shalizi, and M. E. J. Newman (2009): "Power-Law Distributions in Empirical Data," *SIAM Rev.*, 51(4), 661–703.

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.

Cumming, F., and J. Noss (2013): “Financial Stability Paper No 26: Assessing the Adequacy of CCPs’ Default Resources,” Bank of England Financial Stability Papers 26, Bank of England.

Davidoff, T. (2016): “Supply Constraints Are Not Valid Instrumental Variables for Home Prices because They Are Correlated With Many Demand Factors,” *Critical Finance Review*, 5(2), 177–206.

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.

Eaton, J., S. Kortum, and F. Kramarz (2011): “An Anatomy of International Trade: Evidence from French Firms,” *Econometrica*, 79(5), 1453–1498.

Fazio, G., and M. Modica (2012): "Pareto or Log-Normal? A Recursive-Truncation Approach to the Distribution of (All) Cities," Working papers, Business School - Economics, University of Glasgow.

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

Foerster, A. T., P.-D. G. Sarte, and M. W. Watson (2011): "Sectoral Versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production," *Journal of Political Economy*, 119(1), 1–38.

Fogli, A., E. Hill, and F. Perri (2012): "The Geography of the Great Recession," in *NBER International Seminar on Macroeconomics 2012*, NBER Chapters, pp. 305–331. National Bureau of Economic Research, Inc.

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

Fujiwara, Y. (2004): "Zipf Law in Firms Bankruptcy," *Physica A: Statistical Mechanics and its Applications*, 337(1-2), 219–230.

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

Gabaix, X. (2009): "Power Laws in Economics and Finance," *Annual Review of Economics*, 1(1), 255–294.

——— (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79(3), 733–772.

Gabaix, X., and A. Landier (2008): "Why has CEO Pay Increased So Much?," *The Quarterly Journal of Economics*, 123(1), 49–100.

Geerolf, F. (2016): "A Theory of Pareto Distributions," .

Giroud, X., and H. M. Mueller (2017a): "Firm Leverage, Consumer Demand, and Employment Losses During the Great Recession," *The Quarterly Journal of Economics*, 132(1), 271–316.

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

Grassi, B. (2018): "IO in I-O: Size, Industrial Organization, and the Input-Output Network Make a Firm Structurally Important," Working Papers 619, IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University.

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.

Hilber, C. A. L., and W. Vermeulen (2016): "The Impact of Supply Constraints on House Prices in England," *The Economic Journal*, 126(591), 358–405.

Hill, B. M., et al. (1975): "A Simple General Approach to Inference about the Tail of a Distribution," *The annals of statistics*, 3(5), 1163–1174.

Horvath, M. (1998): "Cyclicality and Sectoral Linkages: Aggregate Fluctuations from Independent Sectoral Shocks," *Review of Economic Dynamics*, 1(4), 781–808.

——— (2000): "Sectoral Shocks and Aggregate Fluctuations," *Journal of Monetary Economics*, 45(1), 69 – 106.

Houthakker, H. (1955): "The Pareto Distribution and the Cobb-Douglas Production Function in Activity Analysis," *Review of Economic Studies*, 23(1), 27–31.

Hudson, C., J. Hudson, and B. Morley (2018): "Differing House Price Linkages across UK Regions: A Multi-Dimensional Recursive Ripple Model," *Urban Studies*, 55(8), 1636–1654.

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.

James, O., and P. John (2007): "Public Management at the Ballot Box: Performance Information and Electoral Support for Incumbent English Local Governments," *Journal of Public Administration Research and Theory*, 17(4), 567–580.

Kaplan, G., K. Mitman, and G. L. Violante (2016): "Non-durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data," Working Paper 22232, National Bureau of Economic Research.

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.

Kleiner, K. (2015): "Collateral and Small Firm Labor," Research paper no. 16-49, Kelley School of Business.

Long, J., and C. Plosser (1983): "Real Business Cycles," *Journal of Political Economy*, 91(1), 39–69.

Lucas, R. E. (1977): "Understanding Business Cycles," *Carnegie-Rochester Conference Series on Public Policy*, 5(1), 7–29.

——— (1978): "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 9(2), 508–523.

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

Luckstead, J., and S. Devadoss (2017): "Pareto Tails and Lognormal Body of US Cities Size Distribution," *Physica A: Statistical Mechanics and its Applications*, 465, 573–578.

LUTTMER, E. G. J. (2007): "Selection, Growth, and the Size Distribution of Firms," *The Quarterly Journal of Economics*, 122(3), 1103–1144.

MCLEAN, I., AND A. MCMILLAN (2003): "The Distribution of Public Expenditure across the UK Regions," *Fiscal Studies*, 24(1), 45–71.

MELITZ, M. J., AND G. I. P. OTTAVIANO (2008): "Market Size, Trade, and Productivity," *Review of Economic Studies*, 75(1), 295–316.

MIAN, A., K. RAO, AND A. SUFI (2013): "Household Balance Sheets, Consumption, and the Economic Slump," *The Quarterly Journal of Economics*.

MIAN, A., AND A. SUFI (2011): "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis," *American Economic Review*, 101(5), 2132–56.

——— (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.

Pinter, G. (2015): "House Prices and Job Losses," Discussion Papers 1507, Centre for Macroeconomics (CFM).

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.

Reed, W. J., and M. Jorgensen (2004): "The Double Pareto-Lognormal Distribution - a New Parametric Model for Size Distributions," *Communications in Statistics-Theory and Methods*, 33(8), 1733–1753.

Revelli, F. (2001): "Spatial Patterns in Local Taxation: Tax Mimicking or Error Mimicking?," *Applied Economics*, 33(9), 1101–1107.

Richards, T. (2002): *Was Your Grandfather a Railwayman?: A Directory of Records Relating to Staff Employed by Railways in the Following Countries with Details of Material and Repositories: United Kingdom, Australia, Canada, Eire, India, New Zealand, United States of America and Reference to Crown Agents for the Colonies*. T. Richards.

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.

Stella, A. (2015): “Firm Dynamics and the Origins of Aggregate Fluctuations,” *Journal of Economic Dynamics and Control*, 55, 71 – 88.

Stroebel, J., and J. Vavra (2014): “House Prices, Local Demand, and Retail Prices,” Working Paper 20710, National Bureau of Economic Research.

Tahbaz-Salehi, A., V. Carvalho, M. Nirei, and Y. Saito (2017): “Supply Chain Disruptions: Evidence from the Great East Japan Earthquake,” Discussion paper, Becker Friedman Institute for Research In Economics.

Tesco PLC (2017): “Tesco, 2017 Preliminary Results,” `https://www.tescoplc.com/media/392168/tesco-plc-preliminary-results-1617.pdf`.

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

Vuong, Q. H. (1989): “Likelihood ratio tests for model selection and non-nested hypotheses,” *Econometrica: Journal of the Econometric Society*, pp. 307–333.

Appendix A

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.

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.

A.1 Locally Non-Tradable Firms

Table A.1: 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			

A.2 Locally Tradable Firms

Table A.2: 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			

Appendix B

Mortgages and Interest Rates

B.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.[1] 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

[1] The PSD only provides information on the stock of mortgages after 2012.

very top of the LTI distribution given the evidence that these are most vulnerable to default (e.g., Aron and Muellbauer (2016)).[2] 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.[3] After this process is complete, I am left with around 5.8m mortgages. Figure B.1 provides an estimate of the mortgages that might be missing from the sample.

Figure B.1: Missing Mortgages

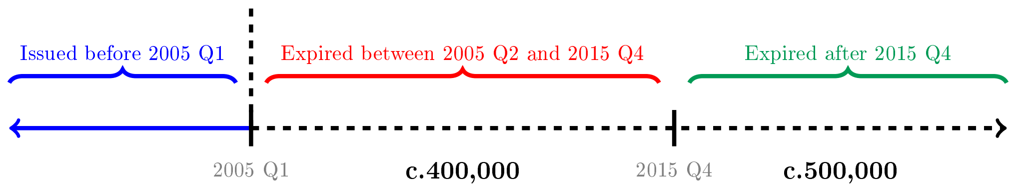

Figure B.1 shows an estimate of the number of mortgages that might be missing

[2]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.

[3]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).

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 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.

B.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.[4]

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:

$$
\begin{aligned}
r_i^{y,type} &= LTV_i + \beta_1 \times deal_i + \beta_2 \times ratetype_i + \beta_3 \times repayment_i \\
&+ \beta_4 \times income_i\{single_i\} + \beta_5 \times income_i\{joint_i\} + \varepsilon_i
\end{aligned} \tag{B.1}
$$

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 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

[4] 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.

at the end of 2008, which means the estimated interest rate levels have a relatively wide tolerance for their accuracy.[5] 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 B.2.

The regressions for first time buyers and remortgagors are reported in Tables B.1 and B.2. 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 B.2. 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.[6] 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

[5] 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.

[6] 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.

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.[7] 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.

Figure B.2: Pre-Crisis LTV Curves

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

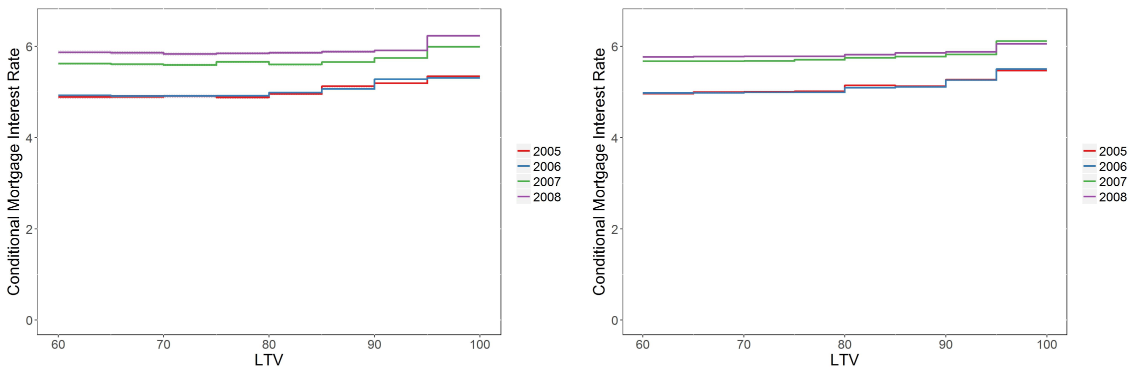

Theses charts use coefficients in Tables B.1 and B.2 to give an indication of mortgage interest rates offered, conditional on mean values for the covariates.

[7]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.

Table B.1: 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 B.2: 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.

Appendix C

Additional Cash-Flow Shock Robustness

C.1 Neighbourhood Sample

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 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 C.1 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 C.1 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 (2017b)), I argue the central specification I use is the most helpful for understanding the cash-flow shock at the micro level.

C.2 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 C.2 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 C.2 confirm the employment-effect estimates are not significant at the 5-percent level.

Table C.1: Sample and Growth Definition Choices

	Annual employment growth		
	(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 C.2: Employment Response Timing

	Annual employment growth		
	(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.

Appendix D

Council Tax Shocks

Local authorities have been able to set a tax on residents since 1991, largely based or property prices, to contribute to the funding of local government. Council tax rates are predominantly based on the historic market value of domestic residences and is paid by the majority of households, though there are various exemptions for single occupancy, vacant possession and low-income residents. Around a quarter of local authorities' revenues comes from council tax, with the remaining budgets made up by central government grants. As these grants tend to be relatively fixed over time, that means council tax is often the only margin of adjustment if local authorities want to increase spending in a given year (see Adam and Browne (2006)).

I use data for 325 English local authorities between 2007 and 2011. In 2011 the Secretary of State for Communities and Local Government announced a blanket freeze in council tax following a consultation on steep and heterogeneous changes since 2001.[1] Although my council tax time series is shorter than for my house price data, I can exploit the significant variation from 2007. Of the 64 missing local authorities, the Scottish parliament worked with the Scottish councils to freeze council tax since 2007 and Northern Ireland operates on a different system of domestic rates, which includes some utilities in the total bill.

Local councils, the administrative body within the local authority, set the tax rate in Great Britain for properties that were worth between £68,001 and £88,000 in April 1991 (so-called Band-D properties).[2] Central government fixes the council tax ratio between Band-D properties and all others bands, which means that in practice local authorities increase council tax for all residents by the same percentage. For example, Band-A properties, which were worth less than £40,000 in 1991, pay two thirds of the Band D tax rate.

Figure D.1: Council Tax Through Time

(a) Nominal Changes **(b)** Percentage Changes

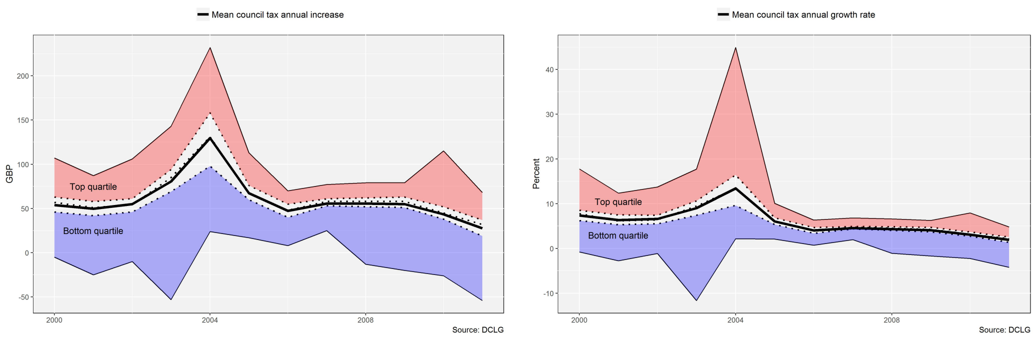

Figures show the swathe of council tax changes across 325 local authorities. After Wandsworth increased its Band-D bill by 44% in 2004 it was still just over half of the average local authority's tax rate.

Figure D.1 plots changes to the council tax charged to the average house in

[1]See https://www.gov.uk/government/speeches/council-tax-freeze-2011-to-2012.

[2]There are other bodies that can levy local taxes on residents and these depend on the exact local devolutionary arrangements.

each English local authority between 2000 and 2011. The first chart in Figure D.1 shows the dispersion of absolute sterling increases in council tax. The black line shows the median local authority. The second chart in Figure D.1 uses the same idea, this time showing how the dispersion in annual percentage change has evolved over time. It does not account for the fact that some of the local authorities with the largest increase may be starting from relatively low bases. But the level of the council tax appears to be unrelated to its growth rate. Between 2000 and 2011 the correlation between absolute value and growth rate for each of the 325 local authorities was around -0.002, which suggests that there is very little short-term mean reversion when future taxes are set.

The median UK house is in Band C and the mean UK household total council tax bill in 2007 was £1,056. There is substantial variation between the levels and growth rates of council taxes between local authorities. In 2007 the average yearly council tax across all English local authorities was around £1,340 for a Band-D property. Within that, the Midlands authority of Newark and Sherwood charged the most at around £1,550 a year, while the London Borough of Westminster charged less than half that. Between 2007 and 2008 the average increase in council tax was 4.3% (i.e. around £50), though a couple of local authorities made small council tax reductions and Leicester City and South Kesteven increased their rates by over 6%.