European Economic Review 144 (2022) 104019

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

Mortgage cash-flows and employment

Fergus Cumming[1]

Bank of England, London, EC2R 8AH, United Kingdom

ARTICLE INFO

JEL classification:
E24
E52
G21

Keywords:
Employment
Interest rates
Monetary policy
Mortgages

ABSTRACT

This paper quantifies the impact of the cash-flow channel of monetary policy on employment by combining novel micro datasets with near-universal coverage. When interest rates fall, mortgaged households spend an important share of the extra cash-flow in their local economy, increasing local labor demand. Overall, a reduction in mortgage payments of £1,000 per household leads to a 0.5 percentage point increase in locally non-tradable employment growth over three years, almost entirely driven by the restaurant sector. Employment-growth effects operate through the intensive margin for small establishments, lower death rates for the mid-sized and higher birth rates for the largest outlets.

1. Introduction

Monetary policy operates through a number of direct channels, but little is understood about the heterogeneity of its transmission through different industries. One channel that has recently attracted attention operates through the shock to household disposable income that results from the interaction of policy interest rates and mortgage debt. This cash-flow effect has been considered by policymakers for some time (e.g., Bernanke (2007)) but it is often difficult to quantify how it might directly affect employment via increases in local consumption and labor demand. I add to our understanding of this topic by combining multiple datasets with near-universal coverage to investigate the importance of the cash-flow effect on employment and firm heterogeneity.

A recent strand of literature has explored the transmission of macroeconomic shocks across the economy in the context of debt, both for households (e.g., Mian and Sufi (2014); Verner and Gyongyosi (2018); Cloyne et al. (2020)) and firms (e.g., Giroud and Mueller (2019) and Gürkaynak et al. (2019)). I exploit regional heterogeneity in the concentration of adjustable-rate UK mortgages resulting from the timing of mortgaging activity to identify the effect of neighborhood cash-flow shocks on locally non-tradable employment during the Great Recession. The 400 basis point reduction in policy interest rates between October 2008 and March 2009 reduced households' repayments at different points over the next few years. Neighborhoods with a large share of adjustable-rate mortgages in 2008 soon received substantial favorable cash-flow shocks, which slackened liquidity constraints and boosted local consumption and employment.

The central analysis is split into two stages. I first model the evolution of around five million mortgage payment flows to estimate the household-level change in cash flows that followed from the systematic and unexpected easing of monetary policy

E-mail address: fergus.cumming@bankofengland.co.uk.

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

https://doi.org/10.1016/j.euroecorev.2021.104019
Received 14 February 2021; Received in revised form 25 July 2021; Accepted 12 September 2021
Available online 26 February 2022

in the fall of 2008. I go on to define locally non-tradable establishments as those generating revenues from nearby customers and group households into neighborhoods.[2] I then exploit the staggered timing of interest rate pass-through to estimate the relationship between changes in cumulative neighborhood-level cash-flow shocks in the seven quarters after the dramatic fall in interest rates, and the growth of locally non-tradable employment between 2007 and 2010.

For many mortgaged households, the direct effect of monetary policy on household finances was substantial. I document that the average household on a fixed-rate mortgage in the middle of 2008 experienced a reduction in mortgage repayments equivalent to £3,700 following the quarters after policy interest rates began their steep descent at the end of 2008. The size of this windfall largely depended on when their mortgage automatically reset to an adjustable rate following the end of their initial fixation period. That was often pre-determined by mortgaging choices in bygone years. In contrast, households on an adjustable-rate mortgage throughout the period of study experienced a total reduction in repayments of around £6300, equivalent to 10.6% of gross annual income. Aggregating these windfalls to around 9000 neighborhoods (so-called *wards* in the UK) provides the variation required to identify the knock-on impact to employment in the local economy.

I find that locally non-tradable establishments in neighborhoods with a large fraction of adjustable-rate mortgagors increased their employment relative to neighborhoods awash with fixed-rate mortgagors. My central estimates suggest that a neighborhood-wide average cash-flow shock of £1,000 per mortgagor over seven quarters was associated with just under a 0.5 percentage point (pp) increase in the annual growth of locally non-tradable employment between 2007 and 2010 in that neighborhood. On average, my results suggest that a total cash-flow injection in a neighborhood of around £125,000 saves one locally non-tradable job over that time horizon, which is consistent with estimates of the marginal propensity to consume out of windfalls (e.g., Bunn et al. (2018); Christelis et al. (2019) and Fagereng et al. (2019)) and salaries in the locally non-tradable sector. The cash-flow effect disappears entirely when an alternative specification is run on manufacturing-sector employment, whose establishments' demand for labor is unlikely to be related to consumption by local residents. Comparing locally non-tradable and manufacturing establishments helps to rule out other confounding channels that might be operating, such as labor-supply effects.

My second contribution is to document *how* that cash-flow boost translated into an overall employment effect. In this regard, I document four facts. First, the cash-flow effect of employment works entirely through the food sector and there is no response in the retail or car-servicing sectors. Second, around a third of the employment effect between 2007 and 2010 operates through the intensive margin of adjustment for the very smallest food-sector establishments — small outlets hiring an additional worker to meet capacity. Entry and exit rates for the food sector are, however, high relative to other industries, so the make-up of establishments in 2010 looked very different to that in 2007. So next, another third of the employment effect is primarily driven by a reduction in the death rate associated with middle-sized establishments employing around 15 people. Fourth, the final third is driven by an increase in the birth rate of the very largest food-sector outlets. Neighborhoods that received positive cash-flow boosts were relatively attractive locations for chains to expand their establishment network. This is the first study to show that employment effects following these types of demand shocks are roughly equally spread across the establishment-size distribution, but operate through different margins of adjustment.

An important concern for the main analysis is that people who took out fixed and adjustable-rate mortgages were different. In particular, my estimates would be biased upwards if the households more likely to take out adjustable-rate mortgages ended up being more insulated from the eventual economic slow-down. I show that observable characteristics across interest rate types were relatively similar. Indeed, since fixation periods were short, the choice between fixed and adjustable was often driven by whichever contract had the lowest initial interest rate, perhaps because people were myopic or liquidity constrained. As a result, almost 40% of those refinancing switched between fixed and adjustable-rate contracts in the immediate run-up to the Global Financial Crisis (henceforth, Crisis).

I propose an instrumental variable approach for my main specification where I construct a *predicted* neighborhood cash-flow shock based primarily on *when* people took out their mortgages in a particular area. Since fixed and adjustable-rate mortgages are priced using information embedded in the yield curve, the yield curve slope shortly before origination was a good predictor of the spread between the benchmark fixed and adjustable-mortgage rates available, and therefore of the type of mortgage households chose. Moreover, short fixes and early repayment fees meant that the majority of transactions in the stock were refinances, and spaced at regular intervals.[3] That meant the exact timing of many transactions was determined by home-purchase choices several years prior. My instrument therefore combines time variation in the slope of the yield curve with the natural variation in when mortgaging activity occurred to strip out any selection bias in mortgage choices across households. When I use this *predicted* neighborhood cash-flow shock to instrument for the cash-flow shocks neighborhoods actually received, the employment-effect estimate is around 25bp lower than the OLS estimate of 0.7pp, indicating some upward bias in the simpler OLS approach.

Another concern is that interest-rate decisions were based on an anticipation of, or reaction to, monetary policy changes. In fact, households were surprised by the large falls in policy rates. Survey evidence shows that only 10% of households in August 2008 expected policy rates to fall substantially in the coming months. And once households had chosen their mortgage contract before the fall of 2008, they were locked in. 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

[2] Locally non-tradable firms and establishments are defined in the spirit of Mian and Sufi (2014), Verner and Gyongyosi (2018), Giroud and Mueller (2019) and Luck and Zimmermann (2020).

[3] The most common fixed-rate mortgage had a two year initial term, so many households were on a two-year refinancing cycle.

interest rate to reset rather than actively seeking out a new contract.[4] This led to a staggered cash-flow effect across households and neighborhoods as interest rates automatically reset to lower levels. I estimate that in the UK at most 7% of people on fixed-rate contracts actively refinanced their mortgages in 2009 based on total refinancing activity during this period. This stands in stark contrast to the US, where refinancing spiked up following the monetary easing (e.g., Beraja et al. (2018)).

Finally, I demonstrate that other channels are not responsible for the regional variation in employment. Parallel employment trends before interest rates fell suggests that results are not driven by heterogeneous local business cycles and I show that the characteristics of neighborhoods above and below the median predicted cash-flow shock were similar. In addition, I employ an extensive set of neighborhood controls that use establishment data to construct proxies for the level and type of economic activity going into the Crisis. I rule out collateral effects by controlling for house prices and house-price changes between the summer of 2008 and the end of 2010. Since businesses predominantly took out variable-rate bank loans before the Crisis, the employment cash-flow effect I find likely operates through local spending rather than the supply side of the economy.[5]

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 et al. (2016); Di Maggio et al. (2017); and Flodén et al. (2020)). I combine this loan-level approach with the more aggregated analysis of Mian and Sufi (2014), Di Maggio et al. (2017) and Verner and Gyongyosi (2018) to examine the employment growth of locally non-tradable establishments in the UK and the evolution of household cash-flows immediately surrounding them. The combination of tightly defined neighborhoods, a plausible instrument and microdata-aggregated controls means I am more confidently able to use cross-sectional variation to infer the causal impact of cash-flow changes on local employment. Improving the quantitative accuracy of such an important elasticity is crucial in order to understand how the economy might respond to monetary policy, or similar consumption shocks, in the future.[6]

My work contributes to the growing literature that establishes 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 et al. (2013) and Cloyne et al. (2020)).[7] But there has recently been renewed attention on the benefits of adjustable-rate contracts, made more important with the observation that changes in disposable income might have large macroeconomic effects (e.g., Violante et al. (2014); Garriga et al. (2017); Piskorski and Seru (2018); Guren et al. (2018) and Greenwald (2018)). One challenge has been to quantify the relative importance of the cash-flow effect in terms of macroeconomic variables.[8] A comparison of the cash-flow-employment effects I estimate and those resulting from household balance sheet deterioration in the Great Recession (e.g., Mian and Sufi (2014) and Verner and Gyongyosi (2018)) suggests that policy reforms that focus on minimizing balance sheet volatility might achieve better overall employment stabilization in future recessions.[9]

Finally, this paper contributes to the literature deepening our understanding of employment responses at a very granular level. Although we have a solid grasp of the importance of firm size (e.g., Moscarini and Postel-Vinay (2012); Fort et al. (2013); Haltiwanger et al. (2013) and Crouzet and Mehrotra (2017)) and the interaction between firm leverage and employment is well established (e.g., Sharpe (1994) and Giroud and Mueller (2017)), there is a relative dearth of studies examining the detailed drivers of employment following macroeconomic shocks. One notable exception is Bahaj et al. (2019), though relative to that work I investigate a more direct channel of monetary policy for different sectors and sizes of establishments.[10] This is important because the exact manifestation of employment changes might help us understand business behavior more broadly (e.g., Aghion et al. (2017)) and have longer term consequences for the economy (e.g., Haltiwanger (2012) or Messer et al. (2016)). This study also bridges the gap to the extensive labor economics literature examining the margins of adjustment in response to wage changes (e.g., Rohlin (2011); Dube et al. (2016); Aaronson et al. (2018); and Luca and Luca (2019)) and crises (e.g., Varum and Rocha (2012)). By isolating the labor demand channel via consumption, my work builds upon these findings and has broader implications for consumption shocks in response to more than just monetary policy changes.

2. Data

2.1. Mortgage data

My analysis uses the universe of residential mortgages issued by UK lenders since April 2005, collected by the Financial Conduct Authority (FCA) and distributed in the Product Sales Database (PSD).[11] It contains a wealth of information on property, borrower

[4] UK mortgage interest rates typically increase with the LTV ratio, which ties collateral values to the cost of refinancing. This was especially true after 2008, see Best et al. (2018) for details.

[5] The majority of business owners lived in different parts of the country to their establishments, so director mortgage windfalls are ruled out by using spatial variation (e.g., Bahaj et al. (2017)). Nevertheless, it is possible that part of the measured cash-flow effect does operate via a cash-flow improvement of small business owners who have borrowed against their primary residence.

[6] See Luck and Zimmermann (2020) for estimates of the employment effect of unconventional policy.

[7] There is also evidence of wider effects of mortgage interest rate pass-through, such as on fertility (Cumming and Dettling, 2020).

[8] 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 Willen (2017)) and outright default (e.g., Aron and Muellbauer (2016); Byrne et al. (2017); and Kartashova and Zhou (2019)) but identification in this area is often restricted to a limited part of the mortgage market.

[9] For example, macroprudential policy (e.g., Aikman et al. (2019)).

[10] In contrast, Chatterjee and Eyigungor (2019) study the more medium-run question of how firms respond to changes in the natural rate of interest.

[11] See https://www.fca.org.uk/data/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. The PSD does not include some informal refinances if they are with the same lender.

and lender characteristics at the time of origination. Using these mortgage flows, I construct an estimate for the stock of mortgages in July 2008, well before the failure of Lehman Brothers and the internationally coordinated policy interventions in the fall that year. Appendix 1.1 goes through the steps needed to transform the mortgage flows into the stock.[12]

The 5 million mortgages in the stock form a large and representative sample of the residential mortgage market in 2008.[13] 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's policy rate (Bank Rate).[14] Often referred to as adjustable-rate mortgages in the US, these are more commonly known as variable-rate mortgages in the UK. Mortgage terms tended to be between 25 and 30 years but the periods governing the path of the interest rate (henceforth, initial period) 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 adjustable or two-year fixed interest rates (and a maximum loan-to-value (LTV) of 75%). 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 initial periods meant that the split between mortgagors on fixed and adjustable rates has always been relatively even.

Following the end of the initial period, 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 refinance interest rates was around 300-400bp (with little variation across lenders), meaning it was usually beneficial for mortgagors to refinance at the end of their initial period (whether on a fixed or adjustable rate). The majority of mortgagors therefore refinanced every two years during the Great Moderation.

One apparent concern with my stock is that the mortgages missing from the sample were in some way different to the others, biasing the results. In fact, because this study uses cross-sectional variation, bias is only likely to arise if the missing mortgages are somehow unevenly distributed across the country. This seems unlikely for two reasons. First, all lenders are captured in the PSD so there are no regions that under-report because of high exposure to a particular institution. Second, mortgages issued before 2005Q2 would have likely been linked to Bank Rate by 2008 because the flow of long-term fixed-rate contracts has always been very low in the UK. Missing mortgages are discussed more in Appendices 1.1 and 1.2. The proportions of adjustable and fixed-rate mortgages in the summer of 2008 are shown in Fig. 1.

I model mortgagor characteristics from the fourth quarter of 2008 through to the second quarter of 2010 and my primary interest is how quarterly 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 transaction date, location, borrower birth date, loan value, property value, household income and how the interest rate contractually varied over time. Other variables have less than complete coverage, often due to heterogeneous reporting practices of the mortgage lenders. Of these, the most important variable left blank is the interest rate at origination. Fortunately, the highly competitive nature of the UK mortgage market means that I can accurately model what the likely interest rate would have been. In addition, the dramatic change in interest rates is more important than the level for cross-sectional variation. Appendix 1.3 addresses how I deal with variables that have incomplete coverage in more detail.

The size of the cumulative cash-flow shock mortgagors received depended on which mortgage they got, and when they got it. 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 an adjustable-rate contract; or beyond the initial period (of either type), and therefore on the SVR. The first group experienced little or no immediate cash-flow shock. The second received a substantial, favorable one. But the third group also often experienced a favorable cash-flow shock because the typical SVR fell in lock-step with mortgage rates at origination. As time went on, more and more households benefited from the positive cash-flow boost. Fig. 2 shows the evolution of Bank Rate, the two most common (new) mortgage contracts, the SVR and the effective interest rate on the stock of residential lending.[15]

2.2. Employment data

The main source of employment data is the Business Structure Database (BSD), which contains 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 establishment level. Since the IDBR is based on Her Majesty's Revenue and Customs tax data, it captures the universe of economically active firms in the UK that are registered for income tax purposes. An enterprise is defined as the smallest combination of legal units that have a degree of autonomy from an enterprise group and can therefore be thought of as the overall business or firm. Enterprises are made up of one or more establishments, or local

[12] From 2015 the FCA began publishing the stock of mortgages, making the process outlined in Appendix 1.1 much more straightforward for future studies. Data quality are increasing over time. Appendix 1.2 compares the 2015 stock data with the flow-methodology used in this paper.

[13] There are no consistent estimates of the true universe of mortgages at this time but contemporary estimates of the stock suggest my sample is likely to represent around 80%–85% of the relevant residential mortgages.

[14] 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.

[15] 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 £1000 in savings. Fig. 2 shows that the SVR fell less than Bank Rate and so the cash-flow shock was larger for mortgagors on adjustable-rate contracts within their contractual maturity. This is taken account of in the cash-flow modeling.

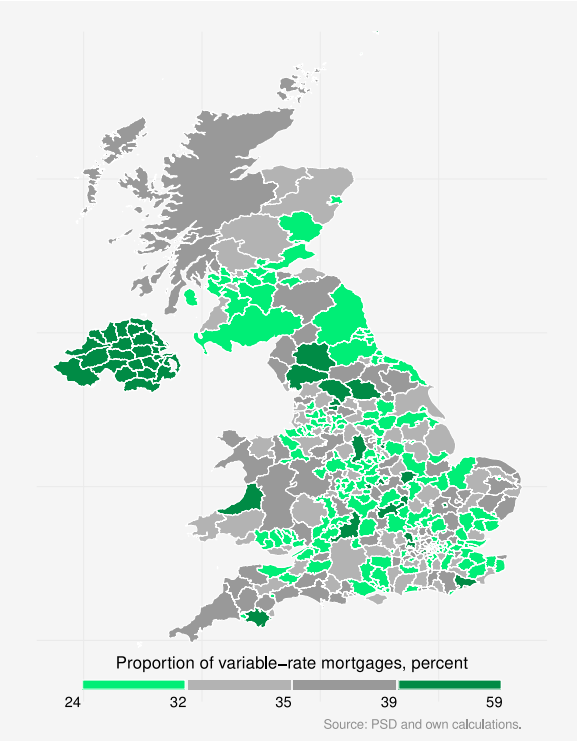

Fig. 1. Distribution of Mortgages in July 2008.
Source: PSD. Figure shows proportion of adjustable-rate mortgages at origination. Color breaks denote quartiles and 389 local authorities are shown.

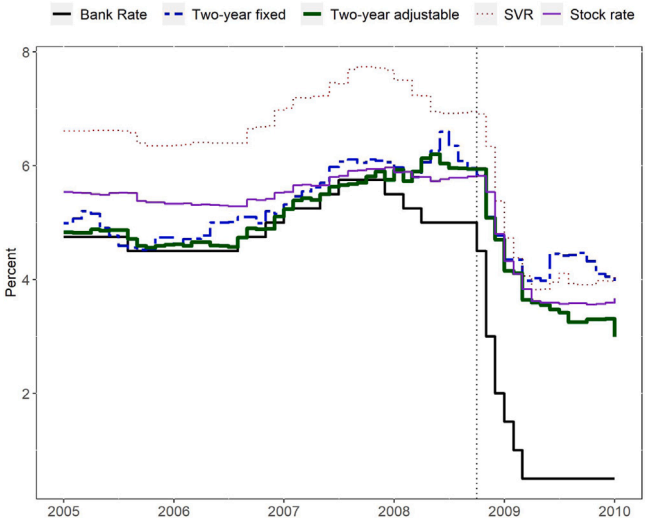

Fig. 2. Mortgage Interest Rates.
Source: Bank of England. Mortgage rates shown (with identifying code): 2-year fixed 75% LTV mortgage (IUMBV34), 2-year adjustable 75% LTV mortgage (IUMBV48), the Standard Variable Rate (IUMTLMV) and the effective interest rate on the stock of lending secured against residential dwellings (CFMHSDE).

units, such as individual shops or restaurants. Businesses are required to report turnover at the enterprise level and employment at the establishment level, as well as geographic information for both.

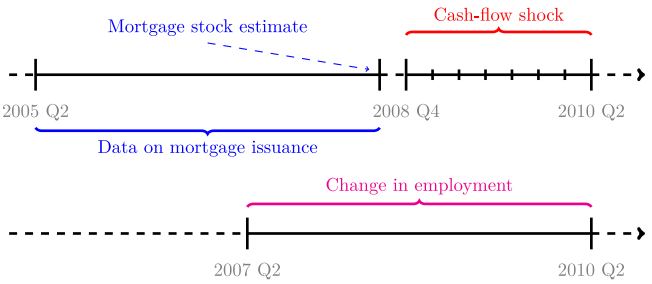

Fig. 3. Empirical Strategy Timeline. The figure shows the timeline for the central specification. I construct an estimate of the stock for the beginning of 2008Q3 using information on mortgages issued since 2005Q2. I then calculate the cumulative cash-flow shock associated with lower mortgage payments for every mortgagor over for the seven quarters from 2008Q4. The dependent variable of the main regression is the percentage change in employment between April 2007 and April 2010. Fig. 11 shows the coefficients for other employment windows using the same cash-flow shock.

To identify the effects of cash-flow shocks, this paper follows Mian and Sufi (2014) and Giroud and Mueller (2019) in categorizing firms into *locally* non-tradable and manufacturing establishments.[16] The original classification for the US is based on a combination of international trade data, geographical concentration as measured by the geographical Herfindahl index and an intuitive sense of which industries respond most to local demand. An equivalent definition in the UK requires a mapping from the four-digit North American Industrial Classification System (NAICS) to the two-digit Standard Industrial Classification (SIC) system used in most of Europe. While the exact mapping between the two systems at different points in time at a very granular level is far from straightforward, the relatively high-level categories used for this analysis in the time period of interest match up almost exactly as can be seen in the first table in Appendix 1.5.[17] The 25 locally non-tradable and 67 manufacturing four-digit NAICS-12 industries therefore map into 13 and 50 SIC-03 groups, respectively, for this study.

Data from the National Transport Survey (see Department for Transport (2016)) suggest that demand for locally non-tradable purchases is relatively tightly defined. Across England, the average shopping excursion is 7 km 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.[18] To reflect these patterns, the baseline specification separates the UK into around 9000 wards, which contain around 5000 people each.

3. Research design

3.1. Baseline strategy

My empirical strategy tests to what extent cash-flow shocks support consumption and its knock-on effect to local employment.[19] Between the fall of 2008 and early spring of 2009 the Bank of England cut its policy interest rate from 4.5% to 0.5%. Such an unprecedented monetary easing mitigated the shock to consumption and employment from the ongoing decline in economic activity and house prices. But even if all households benefited from the support to asset values and their net wealth, only those on adjustable-rate mortgages benefited from an immediate decrease in mortgage payments.[20] This paper exploits regional heterogeneity in the timing of when households received this substantial cash-flow boost. The higher the proportion of households on an adjustable-rate mortgage in the summer of 2008, or transitioning to one soon after rates fell, the more we might expect an increase in spending on local goods and services 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 growth at these establishments to the extent that firms adjust their labor inputs in response to demand shocks.[21]

The heart of this study uses the spatial variation in the timing of the cash-flow shock mortgagors received at the end of 2008 to explain subsequent changes in locally non-tradable employment. Specifically, as shown in the schematic in Fig. 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

[16] On the whole, in the locally non-tradable sector, all establishments belonging to an enterprise are in the same two-digit sic-code. Results are similar when splitting by establishments within a homogeneous and multi-sic-code enterprise.

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

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

[19] By using spatial variation, this work can be thought of as the cash-flow analogue to work such as Nakamura and Steinsson (2014), Dupor and Guerrero (2017) and Verner and Gyongyosi (2018).

[20] Byrne et al. (2017) exploit the difference between SVR mortgages and (policy rate) tracker mortgages for the Irish mortgage market. In the UK the policy interest rate also fell further than the SVR but the difference was less sharp.

[21] 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 Section 6.

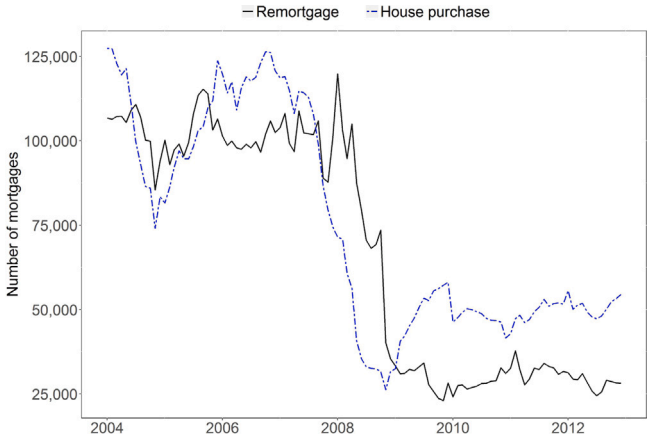

Fig. 4. Mortgage Approvals.
Source: Bank of England. The series correspond to the seasonally adjusted data for housing refinance (B4B3) and housing purchase (VTVX) from the Bank of England Statistical Interactive Database.

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.[22] I then model the loan-level cash flows from the end of 2008Q3 to 2010Q2 and define the cash-flow shock as the change in mortgage payments relative to the counterfactual payments had interest rates not fallen over the intervening quarters.[23] Finally, I compare the neighborhood cash-flow shocks to the three-year change in locally non-tradable employment, between 2007Q2 and 2010Q2.[24]

Previous studies such as Di Maggio et al. (2017) have aggregated mortgage characteristics and changes in employment to relatively small contiguous administrative boundaries such as zip codes or counties. My preferred specification aggregates employment and mortgage characteristics to *wards*, which are a subset of the UK's 389 *local authorities* and contain around 5,000 households. Following Davis et al. (2007), the main outcome variable is e_i^{10-07}, which represents the percentage change in locally non-tradable employment in ward i between April 2007 and April 2010.

$$e_i^{10-07} = 2 \times \frac{E_{i,2010}^{NT} - E_{i,2007}^{NT}}{E_{i,2010}^{NT} + E_{i,2007}^{NT}} \tag{1}$$

After modeling the likely trajectory of every mortgagor's cash flows, I calculate the average cash-flow shock for each neighborhood. In Eq. (2) the household cumulative cash-flow shock is the sum of the quarterly cash-flow shocks, which are defined as the difference between the modeled quarterly mortgage payment following the fall in Bank Rate, p_n^t, and the modeled quarterly payment assuming Bank Rate had not changed, $\tilde{p}_n^t$, for each mortgagor, n, at time, t. The average (sterling-amount) cash-flow shock for neighborhood i is the cumulative cash-flow shock averaged across all mortgagors in the neighborhood, N_i.

$$\overline{\Delta C}_i = \frac{1}{N_i} \sum_{n=1}^{N_i} \sum_{t=2008Q4}^{T=2010Q2} p_n^t - \tilde{p}_n^t \quad \forall\, n \in i \tag{2}$$

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

$$e_i^{10-07} = \alpha + \beta \times \overline{\Delta C}_i + \gamma_j + \delta \times X_i + \varepsilon_i \tag{3}$$

The baseline OLS specification is a regression of employment growth on the average cash-flow shock with fixed effects for local authority j. Identification in this regression comes from variation across neighborhoods within a particular local authority and β yields the percentage point change in locally non-tradable employment growth for each additional £1,000 of mortgage payment reduction per mortgagor. All regressions are weighted by neighborhood employment and standard errors are clustered by the 12 UK regions. Henceforth, I refer to these as my neighborhood regressions.

In some specifications I include a vector of neighborhood-specific controls, X_i, such as the change in house prices and change in local gross-value added (GVA) (excluding wholesale and retail trade).[25] These controls typically use the estimated stock of

[22] In November 2008 that figure had risen to 39%. See question 6 in the Inflation Attitudes Survey at https://www.bankofengland.co.uk/statistics/research-datasets.

[23] This window makes most sense in terms of matching up with the employment window I consider.

[24] To the best of the ONS' knowledge, the vast majority of the annual employment data are recorded shortly after the financial year ends.

[25] Due to data limitations, house price controls are constructed at the local-authority level. The change is house prices and GVA is measured between mid 2008 and 2010Q2.

Table 1
Mortgagor statistics in July 2008.

Statistic	N	Mean	St. dev.	25th pctl	Median	75th pctl
Adjustable-rate						
Age	1,671,530	41.17	10.27	34	40	48
Single income	830,384	57,624	233,487	25,000	36,000	57,500
Joint income	841,146	68,228	131,881	36,000	50,000	72,911
Loan size	1,671,530	141,830	137,097	70,000	110,497	170,599
House price	1,671,530	254,941	241,949	140,000	200,000	293,000
LTV	1,671,530	59.09	24.61	39.30	60.70	80.30
LTI	1,671,530	2.66	1.19	1.86	2.64	3.36
Initial period	1,671,530	2.21	1.21	2	2	2
Fixed-rate						
Age	3,331,885	37.72	10.08	30	37	44
Single income	1,557,485	41,528	110,587	22,500	30,500	45,000
Joint income	1,774,398	53,578	90,507	33,520	43,657	59,177
Loan size	3,331,885	128,298	95,722	75,000	110,000	155,726
House price	3,331,885	202,015	170,507	122,000	167,000	235,000
LTV	3,331,885	67.17	23.54	50.18	72.29	87.88
LTI	3,331,885	2.94	1.10	2.27	2.96	3.60
Initial period	3,331,885	2.48	1.36	2	2	2

This table shows the observable characteristics of mortgages at the end of 2008Q2. Interest-rate types are those designated at origination. Initial periods refer to the length of the contract before the interest rate resets to the Standard Variable Rate.

mortgages and establishment employment data, so accurately capture the characteristics of mortgagors and economic conditions in each neighborhood (see Table 4). The combination of geographic accuracy and near-complete coverage of mortgage data for the regression controls is vital for the causal interpretation of β in this OLS regression.

3.2. Identification challenges

A key identification assumption that spans this research design is that households with adjustable-rate mortgages increased their consumption because they were fortunate enough to receive a cash-flow shock rather than because they differed in other ways to fixed-rate mortgagors. In addition, business employment responses were only driven by changes in relative demand rather than regional heterogeneity in the business cycle.

The first concern regarding identification is that some households might have switched between fixed and adjustable-rate mortgages after the fall in interest rates. This could lead to biased estimates if the propensity to switch mortgages in response to the fall in rates was correlated with mortgagor characteristics. In fact, this active leakage was very small as refinancing activity fell dramatically following the collapse of Lehman Brothers. As collateral values fell at the end of 2008, mortgage lenders withdrew from the once-active refinance market. Even if households were able to refinance in response to falling rates, 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.

Fig. 4 shows how the number of refinances per month fell by around 75% by the start of 2009. Even if all refinances over the next twelve months were fixed-rate mortgagors keen to refinance, that only accounts for around 7% of the stock of mortgages used in the main regressions.[26] This subdued refinancing activity stands in stark contrast to the US where refinances spiked up following the fall in interest rates as those on long term fixed-rate mortgage contracts took advantage of the lower rates. In the UK, the SVR was often an attractive alternative to refinancing. 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.

Although many fixed-rate mortgagors would have seen a positive cash-flow boost on the horizon, a key identification assumption is that households' consumption did not react until the shock arrived. This is based on two pieces of evidence. First, the median mortgagor had less than £1,000 in savings during this period, suggesting that many were unable to use their accumulated resources to smooth their consumption.[27] Second, many mortgaged households were financially constrained from accessing credit. The NMG Survey of Household finances in 2009 showed that almost half of mortgagors and were struggling to keep up with bills and more than half faced an actual or perceived credit constraint.[28] Given the challenging economic environment in 2009 and 2010, it seems plausible that households waited to receive their cash-flow boost before increasing consumption of discretionary items.

Rather than a threat to identification, I can therefore exploit the passive leakage between mortgage types because the exact point at which people rolled on to the adjustable rate was usually determined by previous mortgaging behavior. That meant a neighborhood's seven-quarter cash-flow boost was largely determined by the proportion of adjustable-rate mortgages and the timing of previous mortgaging activity of those on fixed-rate contracts.

[26] There were also around 400,000 home movers during this period.
[27] See Footnote 15.
[28] See Appendix 1.6 for more details.

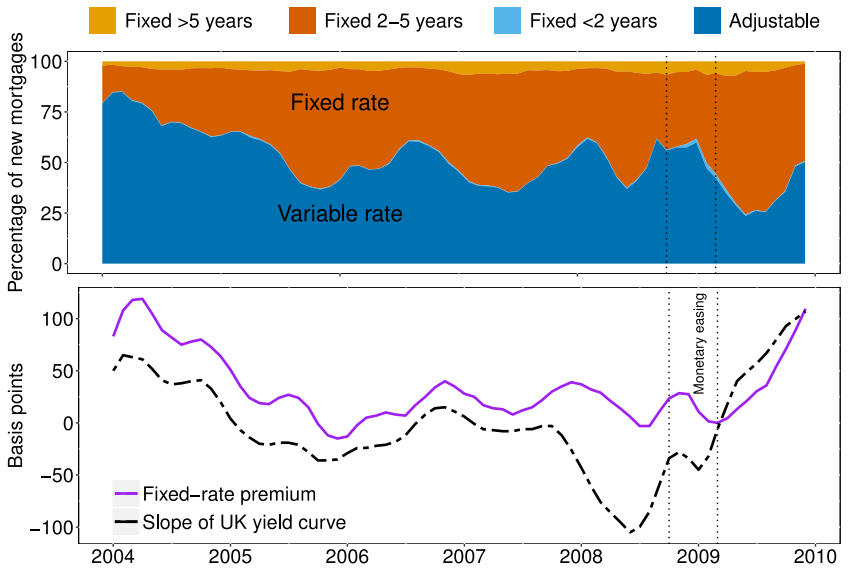

Fig. 5. Mortgage Issuance.
Source: Bank of England. The top panel of this figure shows the variation in mortgage-type issuance across time. These data are taken from Bank of England aggregated data and therefore have a slightly different make-up from the rest of the analysis (e.g., they include second-charge 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 bottom panel shows the time variation in the fixed-rate premium and the slope of the UK yield curve. The former is defined as the difference between the benchmark 2-year 75% LTV fixed-rate mortgage and the 2-year 75% LTV adjustable-rate mortgage. The latter is defined as the difference between 2-year UK Treasury rate and the 0.25-year UK Treasury rate. To allow for delays in the mortgaging process, both series are lagged by three months.

The second major concern is that certain types of households selected into fixed and adjustable-rate mortgages between 2005 and 2008. If more financially secure households opted for an adjustable-rate mortgage this might lead to biased estimates if these households had a lower marginal propensity to consume.[29] Table 1 presents some mortgagor statistics by the type of mortgage chosen an origination. On the whole, the characteristics are broadly similar, though fixed-rate mortgagors tend to be slightly younger and with slightly lower incomes. Although they take out similar sized loans, fixed-rate mortgagors are slightly more leveraged relative to income and collateral values.

Table 1 does not provide enough evidence that Eq. (3) yields unbiased estimates for the employment response to a cash-flow shock. I therefore construct an instrument that strips out selection bias in the choice between fixed and adjustable-rate mortgages at origination.

4. Instrumenting for neighborhood cash-flow shocks

The instrument I construct relies on two key ingredients. First, at least part of households' decisions on whether to pick a fixed or adjustable-rate mortgage was governed by the relative initial price of the two types of contract. Second, at least part of the *exact* timing of the mortgaging decision was determined by factors unrelated to the business cycle. Since the slope of the yield curve moved around between 2005 and 2008, so did the relative price of fixed-rate mortgages, and the fraction of people taking out adjustable-rate mortgages in a given month. I can therefore use the distribution over *when* people got their mortgage in a particular neighborhood to predict the proportion of people on an adjustable-rate contract in the summer of 2008. Moreover, I can use the same distribution to predict how that fraction of people evolved up until the second quarter of 2010 as peoples' fixed-rate contracts expired and they reset onto the SVR. I can then use the neighborhood-level cash-flow shock associated with this predicted adjustable-rate-share evolution to instrument for the actual cash-flow shock the neighborhood received.

[29] To the extent wealthier households are often better able to tolerate financial risk, this would be consistent with the theoretical findings of Campbell and Cocco (2003). That said, the direction of bias is unclear given the recent work on high marginal propensities of consumption for the *wealthy hand-to-mouth* (e.g., Violante et al. (2014)).

Table 2
Evidence for Quasi-Random mortgage selection.

Expected | no intertemporal relationship, percent		
First mortgage	Second mortgage	
	Adjustable	Fixed
Adjustable	13.0	23.0
Fixed	23.0	41.0

Note: Pr(*Fixed*) = 64%.

Observed transition matrix, percent		
First mortgage	Second mortgage	
	Adjustable	Fixed
Adjustable	17.1	18.6
Fixed	18.6	45.7

There were 558,956 people who actively refinanced once between 2005Q2 and 2008Q3 and, in total, 64% of the mortgages were fixed-rate contracts. If the probability of choosing a fixed-rate contract was independent across time one would expect to see the distribution in the top panel. For example, the probability of choosing a fixed-rate contract twice is $0.64 \times 0.64 = 41.0\%$. The observed distribution is shown in the bottom panel.

4.1. Mortgage choice

Mortgages are poorly understood by households and each contract has a vast number of non-price terms and options.[30] But the choice between a fixed and adjustable rate is much less important in the UK than in other parts of the world because the length of most fixed-rate contracts is only a couple of years. Consequently, during the Great Moderation when interest rate volatility was modest, the lifetime cost of choosing an adjustable-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 mortgage choices were often between a 30-year fixed rate with the option to refinance and an adjustable-rate mortgage with a teaser rate.

The top panel of Fig. 5 shows how much variation there was in the origination of fixed and adjustable-rate UK mortgages in the five years before the Crisis. At the start of 2005, for example, two thirds of new mortgages were issued with an adjustable interest rate. Nine months later that proportion had fallen to around a third. I define the fixed-rate premium (FRP) as the spread between the benchmark two-year 75% LTV fixed-rate mortgage and the two-year 75% LTV adjustable-rate mortgage. The UK mortgage market is competitive and there is surprisingly little variation across lenders in the two benchmark rates they offer over time. They also had very little geographical heterogeneity in their lending activities.[31]

The bottom panel of Fig. 5 suggests there was a relationship between the FRP and the likelihood of taking an adjustable-rate mortgage. Coincident with the fall in adjustable-rate mortgage issuance during 2005, the premium associated with the standard fixed-rate mortgage fell by over 50bp. More generally, when the fixed-rate mortgage was more expensive, people were more likely to choose an adjustable-rate contract. The FRP partially captures the market expectation of the future path of UK interest rates and an upward sloping yield curve was associated with people taking out more adjustable-rate mortgages. Before mid 2008, the correlation between the slope of the yield curve and the FRP was 0.77. Fig. 5 therefore suggests that time-varying macro factors played a significant role in determining the types of mortgage people picked.

This mechanism is consistent with evidence from Badarinza et al. (2017) that shows people pay more attention to the headline initial rate than whether the overall cost of the mortgage is good value relative to their expectations about the path of policy rates. Many people struggle to understand the mechanics of mortgage finance and the Bank of England's *Public Attitudes to Inflation* survey shows that people do not consider the market curve when deciding how interest rates might move over the next few quarters.[32] Even if households understood 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.[33] In either case, I can exploit the explanatory power of the yield curve slope in determining mortgage choices as long as mortgagors based at least part of their decision on whether to take out a fixed or adjustable-rate mortgage on the relative attractiveness of the initial interest rate.

Table 2 provides further evidence that decisions were often driven by the slope of the yield curve by using the population of those refinancing. The top panel shows the expected distribution of mortgaging choices under the null hypothesis that the probability of

[30] Agarwal et al. (2015) shows people struggle to calculate the present value of fees compared to additional percentage points on the underlying interest rate, for example.

[31] See Appendix 1.4.

[32] See https://www.bankofengland.co.uk/inflation-attitudes-survey/2010/august-2010. Low mathematical ability is associated with poor financial decision making (Agarwal and Mazumder, 2013).

[33] See Cloyne et al. (2019) for evidence in the UK.

choosing a fixed-rate mortgage was equal to the sample average and independent across time. The bottom panel shows the observed behavior of those refinancing. It shows that over 37% of households that refinanced between 2005 and 2008 moved onto the opposite interest-rate contract, only 9pp lower than expected under the null hypothesis that the characteristics of fixed and adjustable-rate mortgagors are identical. Not only do a surprisingly large share of households flip mortgage type but the switch is symmetric, which is evidence against life-cycle drivers of the choice of interest-rate type.[34] This supports the claim that many households chose their mortgage type based on marginal factors such as the initial cost rather than fixed and adjustable-rate mortgages being segmented markets.

4.2. Predicted cash-flow shock construction

I construct a neighborhood-level cash-flow shock that is primarily based on *when* people in that neighborhood took out their mortgages. It therefore combines exogenous variation in the slope of the yield curve with the natural variation in when people took out mortgages in the neighborhood two years before Bank Rate fell. As before, this is divided by the total number of mortgagors in the neighborhood to give a per-household cumulative sterling amount. This predicted cash-flow shock is used to instrument for the actual cash-flow shock that neighborhoods received.

In order to calculate the neighborhood-level cash-flow shock for Eq. (2), I model the quarterly payments of each of the 5 m mortgages at the end of 2008Q3.[35] This requires taking estimates of the loan amount, loan maturity and time remaining on the fixation period and projecting them forward for the following seven quarters. The cumulative difference between the mortgage payments in this modeling exercise and the counterfactual scenario where Bank Rate had remained constant is the modeled cash-flow shock.

The first step in constructing the instrument is to create an expected cash-flow shock for each household based on their transaction date. This involves taking raw averages of the proportion of adjustable-rate mortgages, the term and the interest rate at origination for each month between April 2005 and June 2008 (denoting them Adj_n^t, $Term_n$ and $Rate_n$, respectively).

These values are reported in Table 3 where the row is the month of origination and the column is the probability of being on an adjustable rate in that quarter. Moving left to right, the values are weakly increasing as more fixed-rate mortgages reset to the SVR. On the whole, moving top to bottom, the values fall as the proportion of households passed their initial period decreases. But the fall is not monotonic because the proportion of newly-originated adjustable-rate mortgages varied with the yield curve slope.

The constructed quarterly payment for a mortgage originated in month m for household h in quarter t uses a simple amortization formula. It combines the loan balance of the household with the average interest rate and number of payments remaining conditional on the month of origination. Subsequent quarterly payments are calculated iteratively and I assume loans are not repaid in the very short term. The key difference between the payment for the adjustable and the fixed versions is that the former interest rate is assumed to move in lockstep with Bank Rate, meaning that quarterly payments fall sharply in 2009. The payment formula is shown in Eq. (4) and the constructed cash flow is shown in Eq. (5).

$$p_{h,v}^t = \frac{Balance_{h,t} \times \widehat{Rate}_{m,v}}{1-(1+\widehat{Rate}_{m,v})^{\widehat{Term}_{m,v}}}, \quad \forall v \in \{Adjustable, Fixed\} \tag{4}$$

$$\Delta C_h = \sum_{t=2008Q4}^{T=2010Q2} p_{h,Fix}^t - p_{h,Adj}^t \tag{5}$$

The next step requires calculating the mean expected cash-flow shock for neighborhood i and month of origination m, shown in Eq. (6). This is done by weighting the household cash-flow shocks estimated in Eq. (5) by the binary probabilities in Table 3, averaging across households. This yields a 39 by 1 vector for each neighborhood (i.e. a row for each month between April 2005 and June 2008).

$$\overline{\Delta C}_i^m = \frac{1}{H_i^m} \sum_{h=1}^{H_i^m} \widehat{Adj}_m \times \Delta C_h \tag{6}$$

To get the predicted cash-flow shock for each neighborhood, this vector needs to be combined with a (39 by 1) weighting vector, θ_i^m, which is shown in Eq. (7). θ_i^m captures the temporal distribution of when mortgages in a particular neighborhood were originated. Although some variation in the predicted cash-flow shock across neighborhoods comes from $\overline{\Delta C}_i^m$, the bulk come from θ_i^m because this determines when fixed-rate mortgagors reset to adjustable rates.

$$\overline{\Delta C}_i = \sum_{m=1}^{39} \theta_i^m \overline{\Delta C}_i^m, \quad where \sum \theta_i^m = 1 \tag{7}$$

[34] For example, it is not typically the case that younger households tend to initially get fixed-rate mortgage before switching to an adjustable-rate mortgage later on.

[35] Given the frequency at which people refinanced, first-time buyers and movers make up a small fraction of transactions. Results are very similar if I restrict my attention purely to those refinancing.

Table 3
Average values for instrument construction.

Month of origination	Probability of being on a variable rate in this quarter										Overall proportion(%)
	2008Q3	2008Q4	2009Q1	2009Q2	2009Q3	2009Q4	2010Q1	2010Q2	Mortgage term	Initial Rate	
2005 M4	0.81	0.81	0.81	0.81	0.82	0.82	0.82	0.82	20.45	5.18	1.33
2005 M5	0.79	0.79	0.79	0.79	0.8	0.8	0.8	0.8	20.64	5.19	1.41
2005 M6	0.79	0.79	0.79	0.79	0.8	0.8	0.8	0.8	20.87	5.2	1.65
2005 M7	0.54	0.75	0.75	0.75	0.75	0.76	0.76	0.76	20.88	5.18	1.71
2005 M8	0.5	0.72	0.72	0.72	0.72	0.73	0.73	0.73	20.87	5.12	1.88
2005 M9	0.45	0.7	0.7	0.7	0.7	0.71	0.71	0.71	20.76	5.07	2.00
2005 M10	0.46	0.46	0.71	0.71	0.72	0.72	0.73	0.73	20.85	5.04	1.97
2005 M11	0.46	0.46	0.71	0.71	0.72	0.72	0.73	0.73	20.91	5.03	2.10
2005 M12	0.46	0.46	0.72	0.72	0.72	0.72	0.73	0.73	21.2	5.04	2.00
2006 M1	0.53	0.53	0.53	0.75	0.75	0.75	0.75	0.76	20.82	5.06	1.76
2006 M2	0.55	0.55	0.55	0.76	0.76	0.76	0.76	0.77	20.98	5.06	1.81
2006 M3	0.57	0.57	0.57	0.76	0.77	0.77	0.77	0.78	21.09	5.07	2.30
2006 M4	0.54	0.54	0.54	0.54	0.74	0.74	0.74	0.74	21.32	5.06	2.14
2006 M5	0.54	0.54	0.54	0.54	0.74	0.74	0.74	0.74	21.31	5.06	2.69
2006 M6	0.57	0.57	0.57	0.57	0.75	0.75	0.75	0.75	21.61	5.07	3.01
2006 M7	0.41	0.68	0.68	0.68	0.68	0.81	0.81	0.81	21.61	5.08	3.00
2006 M8	0.45	0.7	0.7	0.7	0.71	0.82	0.82	0.82	21.74	5.14	3.28
2006 M9	0.45	0.71	0.71	0.71	0.71	0.83	0.83	0.83	21.76	5.18	2.96
2006 M10	0.44	0.44	0.79	0.79	0.79	0.79	0.87	0.87	21.76	5.19	3.21
2006 M11	0.42	0.42	0.79	0.79	0.79	0.79	0.87	0.87	21.69	5.22	3.52
2006 M12	0.39	0.39	0.78	0.78	0.79	0.79	0.87	0.87	22	5.24	3.01
2007 M1	0.34	0.34	0.34	0.79	0.79	0.79	0.79	0.87	21.58	5.49	2.78
2007 M2	0.29	0.29	0.29	0.77	0.77	0.77	0.77	0.86	21.65	5.54	2.60
2007 M3	0.27	0.27	0.27	0.76	0.76	0.76	0.76	0.85	21.74	5.58	3.25
2007 M4	0.27	0.27	0.27	0.27	0.71	0.71	0.71	0.71	21.95	5.62	2.94
2007 M5	0.27	0.27	0.27	0.27	0.71	0.71	0.71	0.71	22.02	5.68	3.31
2007 M6	0.26	0.26	0.26	0.26	0.72	0.72	0.72	0.72	21.95	5.71	3.45
2007 M7	0.23	0.26	0.26	0.26	0.26	0.73	0.73	0.73	21.97	5.78	3.48
2007 M8	0.25	0.29	0.29	0.29	0.29	0.74	0.74	0.74	22.2	5.85	3.49
2007 M9	0.3	0.33	0.33	0.33	0.33	0.76	0.76	0.76	22.03	5.91	3.06
2007 M10	0.32	0.32	0.35	0.35	0.36	0.36	0.84	0.84	21.88	5.94	3.28
2007 M11	0.35	0.35	0.38	0.38	0.38	0.38	0.85	0.85	21.92	5.95	2.93
2007 M12	0.37	0.37	0.4	0.4	0.4	0.4	0.86	0.86	21.97	5.92	2.23
2008 M1	0.43	0.43	0.43	0.46	0.46	0.46	0.46	0.84	21.21	5.9	2.48
2008 M2	0.47	0.47	0.47	0.5	0.5	0.5	0.5	0.83	21.12	5.86	2.42
2008 M3	0.45	0.45	0.45	0.47	0.47	0.47	0.47	0.79	21.32	5.84	2.24
2008 M4	0.4	0.4	0.4	0.4	0.41	0.41	0.41	0.41	21.23	5.79	2.59
2008 M5	0.33	0.33	0.33	0.33	0.34	0.34	0.34	0.34	21.61	5.81	2.43
2008 M6	0.31	0.31	0.31	0.31	0.31	0.31	0.31	0.31	21.52	5.86	2.30
R^2	0.07	0.11	0.15	0.17	0.14	0.11	0.10	0.09	0.00	0.38	

This table summarizes the probability that a refinanced mortgage originated in the row-month is an adjustable-rate mortgage by the time of the column-quarter for 2.6 m refinanced mortgages originated between 2005M4 and 2006M6. Moving left to right, the values are weakly increasing as more fixed-rate mortgages reset to the SVR. On the whole, moving top to bottom, the values fall as the proportion passed their initial period decreases. But the fall is not monotonic, because the proportion of newly originated adjustable-rate mortgages varied with the yield curve slope. Almost all coefficients are signification at the 0.1% level.

The simplest course is to weight the cash-flow shock vector by the distribution of the number of mortgages issued in each month for that neighborhood, which is the approach used in the main specification for simplicity.[36] But in the robustness tests I weight the mean predicted cash-flow shocks by the time distribution of mortgage transactions from Land Registry data two years prior. The distribution of transactions two years prior is correlated with the contemporaneous timing distribution because so many people were on contracts with a two year initial term. Short fixation periods and large spreads to SVR-reset rates meant that people had a strong incentive to refinance when their mortgage deal expired before the Crisis. Over half of the mortgages in the fall of 2008 were refinances. At the same time, hefty early repayment fees on the order of percentage points of loan-value-outstanding limited the attractiveness of applying for to refinance too early. Using matched refinanced mortgages contained within my stock, I estimate that around 40% of people refinanced within a few weeks of the end of their contractual maturity and around three quarters within a few months (see Fig. 6, which is similar to Figure A.1 in Best et al. (2015)). Using this alternative weighting distribution helps satisfy the exclusion restriction in case contemporaneous mortgaging choices are correlated with local economic conditions.

The left-hand panel of Fig. 7 shows the relationship between the predicted cash-flow shock and the actual cash-flow shock per mortgagor in each neighborhood, split into London and non-London wards. It shows that London neighborhoods tended to have larger cash-flow shocks (because property prices were higher) but the instrument relationship was similar. The right-hand panel

[36] Figure A1 in the Appendix illustrates the variation across months-of-origination in the PSD data. As described below, any residual variation from the cash-flow shock vector is dealt with in the main specification by using local-authority fixed effects and adding house price controls.

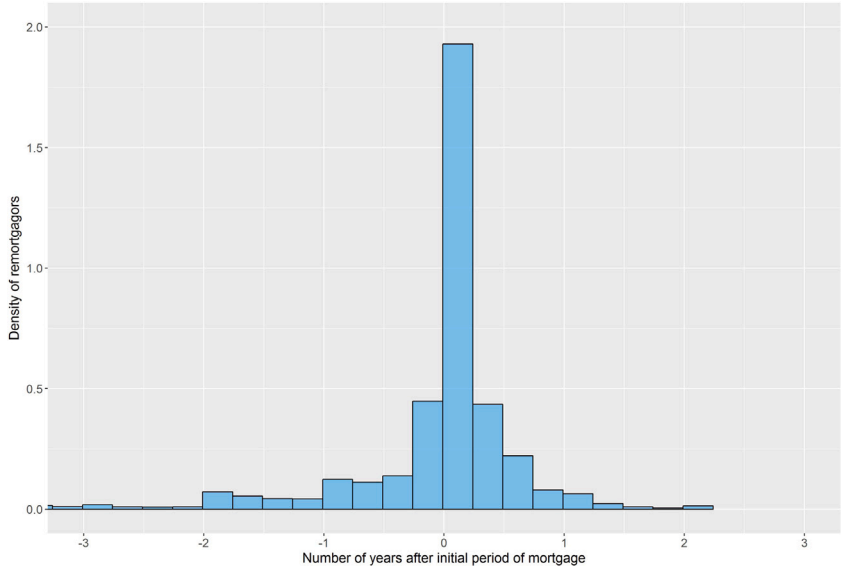

Fig. 6. Timing of Refinancing Activity.
Source: PSD. This figure shows at what point mortgagors chose to refinance their contract using matched refinanced mortgages in my stock. For example, someone who took out a two-year deal and refinanced 25 months later would appear in the modal bucket

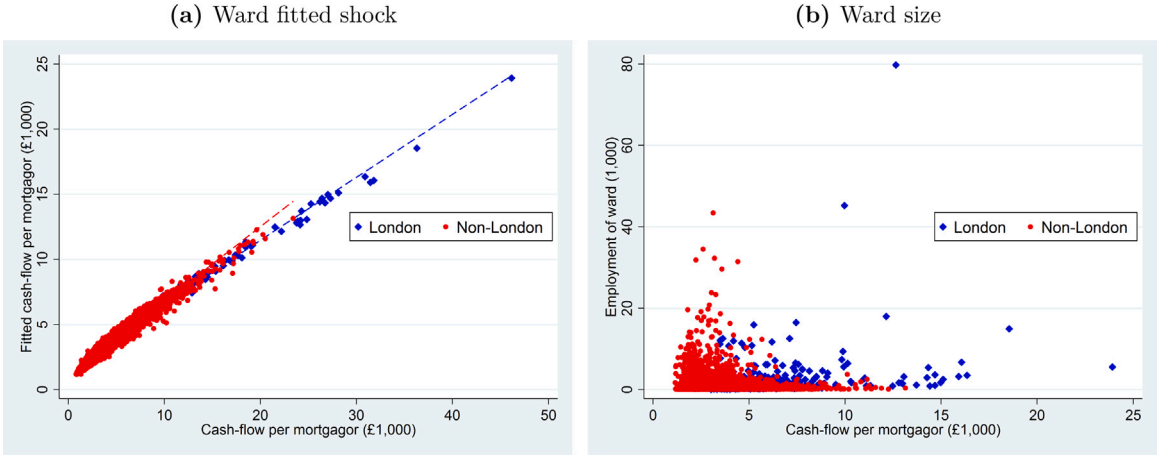

Fig. 7. Neighborhood Instrument.
Source: PSD and ONS. The left-hand panel shows a scatter plot of the average cash-flow shock per mortgagor (2008Q4–2010Q2) versus the predicted cash-flow shock and ward size for c.9000 wards. Dashed lines show instrument fit for London and Non-London wards. On the right-hand panel employment is the average locally non-tradable employment in each ward between 2007 and 2010.

shows the relationship between the cash-flow shock and locally non-tradable employment in each neighborhood. The implication of these figures for the estimated coefficients is explored further in the robustness checks.[37]

Conceptually, variation for the instrument comes from within-local-authority patterns of purchase and refinance. Imagine a local authority with two neighborhoods. The first contains very old housing stock and there was a constant flow of transactions between 2005 and 2008. The second is formed of a new development built at the turn of the century. In the second neighborhood many of the refinancing cycles will be aligned and so there will be a concentration of refinances at some point every two years. Exactly when this falls in the 2005–2008 period will partially determine the relative abundance of adjustable-rate mortgages when policy rates fall in 2008.

[37] Table A8 shows some correlations between the predicted cash-flow shock and neighborhood characteristics.

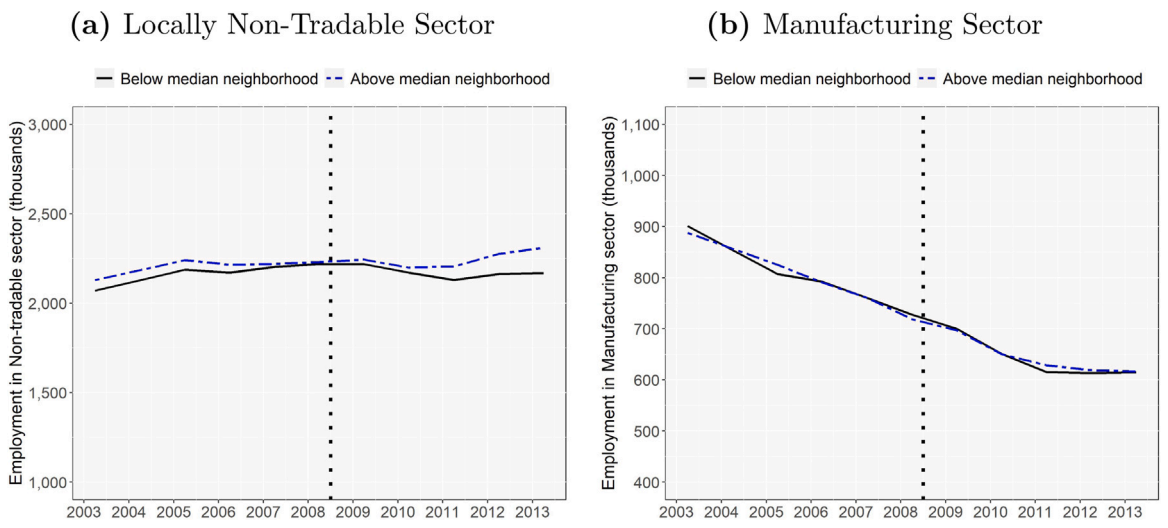

Fig. 8. Employment Trends and Predicted Cash-flow Exposure.
Source: PSD and BSD. These figures show total employment in the locally non-tradable and manufacturing sectors split by wards above (blue dashed) and below (black solid) the median predicted ward cash-flow shock in 2007. The dotted line indicates where mortgage stock is taken.

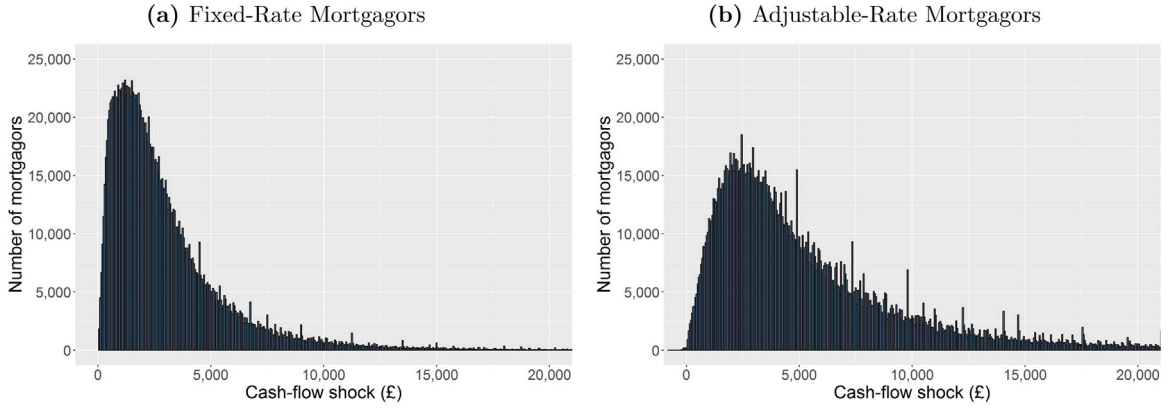

Fig. 9. Individual Cash-flow Shock Distribution.
Source: PSD. Around 5 million mortgages are shown across the two panels. Fixed-rate mortgagors are those on a fixed-rate contract as of 2008Q3 and does not indicate status at origination. Many fixed-rate mortgagors rolled on to the SVR at some point before 2010. Buckets are £50 wide, though the £0–50 bucket has been excluded in the left-hand panel to make visual comparisons easier.

Finally, Fig. 8 shows the employment trends when neighborhoods are categorized as above and below the median predicted neighborhood cash-flow shock. There are three features of this Figure that make identification using the instrument convincing. First, the left-hand panel shows that employment trends in the run-up to the Crisis were roughly parallel, suggesting these neighborhoods were relatively similar apart from the mortgage composition. Second, the slight convergence of the black and blue dotted lines in the left-hand panel demonstrates the sensitivity of the neighborhoods to interest rates: Fig. 2 shows that interest rates increased over 100bp between the middle of 2006 and 2007, which would be expected to depress employment more and after the middle of 2008.

5. Results

5.1. Summary statistics

The timing of mortgage originations and variation in the choice of interest rate type diffracted the direct effect of monetary policy into a heterogeneous cash-flow shock across mortgagors. Fig. 9 shows the distribution of that cash-flow shock across mortgagors on

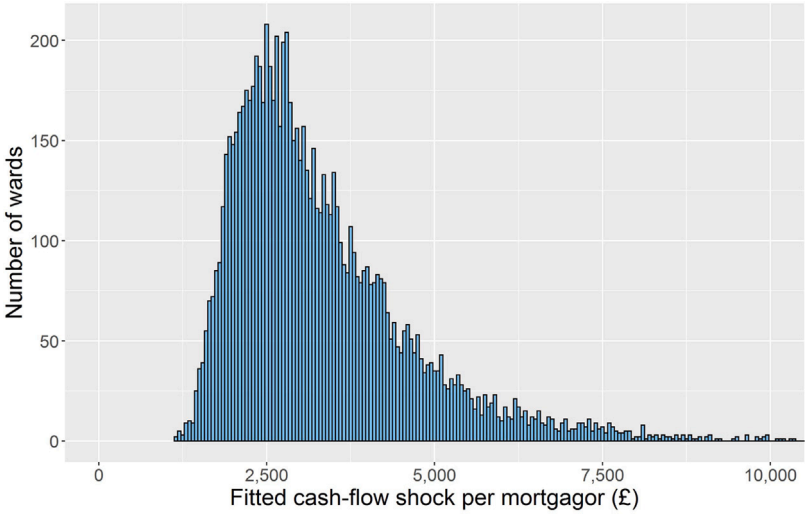

Fig. 10. Neighborhood Cash-flow Shock Distribution.
Source: PSD. This figure plots the mean predicted cash-flow shock across around 9000 wards. The median value is around £3000 and the mean is around £3300.

both fixed and adjustable-rate contracts as of the middle of 2008. The left-hand panel shows that cash-flow shocks for those on fixed-rate mortgages were clustered closer to zero; around a third of these households received no cash-flow shock at all. Nevertheless, because the average length of the fixed period was a little over two years, many fixed-rate households were able to roll onto the SVR at some point and gain some direct benefit of the monetary stimulus before the end of 2010. The mean (median) seven-quarter cash-flow shock for these mortgagors was around £3,700 (£2,200), which was equivalent to around 6.6% (5.6%) of their annual gross income over this period.[38]

Those on adjustable-rate contracts received a mean (median) cash-flow shock of around £6,300 (£4,200), equivalent to around 10.6% (9.3%) of their gross income. This is an economically significant number; similar to the household saving rate during the Great Moderation. For many, it was close to 30% of annual discretionary income (after subtracting essential food, travel, etc.).[39] The spatial distribution of fixed and adjustable-rate mortgagors at different points in time led to heterogeneity in the average cash-flow shock for each neighborhood. Fig. 10 shows the predicted cash-flow shock for each neighborhood when the timing of each mortgaging decision is the sole determinant of the time left on a fixed rate from 2008Q3, described above. This is used to instrument for the actual cash-flow shock in the main regressions.

Table 4 presents the summary statistics associated with the main set of regressions, split by wards above and below the median predicted cash-flow shock. It shows that the main sample consists of over 9000 contiguous neighborhoods, each containing an average of just over 500 mortgages. Around a third of households were on an adjustable-rate contract at origination. The average neighborhood mortgagor cash-flow shock was just under £4,000. The mean neighborhood employed around 490 people in the locally non-tradable sector in 2007, and this constituted around a sixth of overall private-sector employment. The average neighborhood experienced a 1.2% reduction in locally non-tradable employment between 2007 and 2010, though this masks a fair degree of heterogeneity across neighborhoods.

5.2. Local employment cash-flow effect

Table 5 presents the main results of estimating β in Eq. (3) using the definitions in Eqs. (1) and (2). The first column shows the first stage of the IV estimation using the predicted cash-flow shock as an instrument for the actual neighborhood cash-flow shock. The statistically significant coefficient shows that the instrument has high relevance. The second column uses a standard OLS approach and suggests that a cash-flow shock equivalent to £1000 per mortgagor was associated with an increase in locally non-tradable employment growth of 0.7pp and is statistically significant at the 0.1 percent level when standard errors are clustered by the 12 regions. The third column shows the main IV estimates and suggests using the predicted cash-flow shock is important, as it reduces the estimated coefficient to 0.44pp.[40] This provides evidences that using a standard OLS approach produces some upward

[38] If there is a lag between the increase in cash flow and non-tradable employment, the seven-quarter measure of the cash-flow shock for this group might be an underestimate.

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

[40] When clustering at the local-authority level the t value falls to 2.30. When running on the subset of those refinancing, the coefficient is 0.456 and the t value is 4.00.

Table 4
Neighborhood summary statistics.

Statistic	N	Mean	St. dev.	25th pctl	Median	75th pctl
Wards below median fitted cash-flow shock						
Actual cash-flow shock, £	4,662	2.37	0.63	1.88	2.34	2.81
Fitted cash-flow shock, £	4,662	2.34	0.40	2.04	2.37	2.66
Share of flexible-rate mortgages, %	4,662	39.81	7.77	34.50	38.40	43.43
Number of mortgaged households	4,662	587.8	428.9	286	467.5	793
Population	4,662	6,528	4,466	3,343	5,107	9,021
LTI	4,662	2.76	0.18	2.65	2.77	2.89
House price change, 2010Q2-08Q3, %	4,662	−5.78	5.24	−7.86	−5.09	−2.50
Share LNT employment, %	4,662	24.26	14.13	13.41	21.59	32.26
GVA level, £000 per capita	4,662	49.56	13.75	42.53	47.53	53.63
GVA change, 2010–2007, %	4,662	−4.21	4.46	−6.92	−4.32	−1.39
LNT employment 2007	4,662	518.2	920.4	101.2	246	576
LNT employment 2010	4,662	510.4	869.9	97	238	584
Wards above median fitted cash-flow shock						
Actual cash-flow shock, £	4,406	5.61	2.84	3.92	4.81	6.35
Fitted cash-flow shock, £	4,406	4.36	1.48	3.39	3.94	4.81
Share of flexible-rate mortgages, %	4,406	47.67	7.18	42.52	47.42	52.31
Number of mortgaged households	4,406	500.1	369.8	214	401	693
Population	4,406	5,483	3,976	2,328	4,308	7,591
LTI	4,406	2.91	0.17	2.80	2.91	3.03
House price change, 2010Q2-08Q3, %	4,406	−2.38	3.90	−4.17	−2.14	−0.06
Share LNT employment, %	4,656	19.32	11.51	10.89	16.89	25.47
GVA level, £000 per capita	4,656	54.69	18.17	43.74	51.19	61.10
GVA change, 2010–2007, %	4,656	−1.83	3.84	−4.16	−1.73	0.58
LNT employment 2007	4,406	455.3	1,156.2	92	190	430
LNT employment 2010	4,406	451.4	1,160.9	90	189	429.8

This table presents summary statistics for the neighborhood data used in the analysis. The population estimates are constructed using postcode-level census data. House price changes are from the ONS local-authority series. Locally non-tradable (LNT) employment data are derived from the BSD and the neighborhood GVA shocks are constructed using the neighborhood share of employment in each of the 17 main industrial sector categories and regional GVA indices. All other data are constructed from the PSD.

Table 5
Main neighborhood regressions.

		Employment growth, pp			
	First stage	No instrument	Main specification	Di Maggio et al. (AER 2017)	Di Maggio et al. instrumented
Fitted cash-flow (£000)	1.947***				
	(26.09)				
Cash-flow shock (£000)		0.705***	0.444***		
		(7.41)	(5.38)		
Adjustable-rate share (% of mortages)				0.196	0.278**
				(1.95)	(3.02)
Geographic fixed effects	Yes	Yes	Yes	Yes	Yes
Specification	OLS	OLS	IV	OLS	IV
Observations	9,068	9,068	9,068	9,068	9,068
R^2	0.993	0.092	0.005	0.009	0.002
Chi^2			28.95		9.13
F	11,390				

Note: *p<0.05; **p<0.01; ***p<0.001.
The dependent variable is the neighborhood-level locally non-tradable employment growth between 2007 and 2010. Regressions are weighted by employment and standard errors are clustered at the region level (t-statistics shown in parentheses). With the exception of the fourth column, fixed effects are at the local-authority level. The DiMaggio et al. (2017) OLS regression is performed with regional fixed effects. There are 10 regions and 389 local authorities.

bias because of selection into different types of mortgage (e.g., Table 1 suggests households choosing adjustable-rate mortgages tended to have slightly higher incomes).

One of the contributions of this paper is to quantify the employment effect of cash-flow shocks via local spending. There is currently very limited evidence on how large this effect is, though Di Maggio et al. (2017) estimate that every 10pp increase in the zip-code share of adjustable-rate mortgages increases the overall employment growth rate by 0.29pp. Column 4 runs a regression in a similar style where the independent variable is the proportion of adjustable-rate mortgagors in 2008Q3. This estimate suggests that every 10pp increase in adjustable-rate mortgages increases locally non-tradable employment by 2pp, though it is imprecisely estimated. Using the instrument, the fifth column shows that the equivalent coefficient is around 2.8pp for every 10pp of adjustable-rate mortgages. Since locally non-tradable employment represents around a sixth of overall private-sector employment,

Table 6
Robustness.

	Employment growth, pp						
	Standard growth	Log growth	Lagged instr.	LTI interaction	Macro controls	HP Channel	Excl. London
Cash-flow shock	0.474***	0.445***	0.509***	0.580***	0.519***	0.661***	0.409***
(£000)	(5.163)	(5.275)	(8.748)	(8.286)	(7.377)	(7.228)	(2.663)
LTI				−6.338*			
(At origination)				(−2.136)			
High LTI				−0.161			
(Binary interaction)				(−1.469)			
GVA per capita					−0.026	0.007	−0.016
(£000)					(−1.532)	(0.268)	(−0.631)
Manufacturing share					−0.001	−0.001	−0.001
(% of employment)					(−1.244)	(−1.353)	(−1.221)
Employment, 2010–2007					−0.007	−0.007	−0.010
(% change)					(−0.188)	(−0.214)	(−0.284)
House price						0.163**	0.207**
(Index, 2008Q3)						(2.900)	(2.581)
House price change						0.306**	0.525**
(2010Q1–2008Q3, %)						(2.573)	(2.657)
Local-authority fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Specification	IV	IV	IV	IV	IV	IV	IV
Observations	9,068	9,068	8,091	9,068	4,936	4,936	4,790
R^2	0.002	0.004	0.007	0.009	0.009	0.012	0.008

Note: *p<0.05; **p<0.01; ***p<0.001.
Regressions are weighted by employment and standard errors are clustered at the region level (*t*-statistics shown in parentheses). There are 10 regions and 389 local authorities. Standard growth is defined as employment in 2010 divided by employment in 2007. Log growth is defined as the log change in employment. High LTI is defined as neighborhoods above the median LTI. House price changes are from the ONS local-authority series. London is defined as the following local authorities: City, Camden, Greenwich, Hackney, Hammersmith and Fulham, Islington, Kensington and Chelsea, Lambeth, Lewisham, Southwark, Tower Hamlets, Wandsworth and Westminster.

that yields an estimate for total employment growth of around 0.46pp, somewhat higher than the Di Maggio et al. (2017) result, though direct comparisons are made difficult by the differences in employment-growth windows chosen. One plausible reason for any slight discrepancy is that the cross-county US regressions might pick up a correlation between mortgage type and household characteristics.[41] Another possibility is that interest rate transmission through the mortgage market in the UK is stronger than in the US because of differences in the structure of the retail or labor markets.

5.3. Robustness

I perform a number of robustness checks to test the empirical specification and rule out competing channels in Table 6. In the first two columns I investigate whether using alternative definitions for the employment growth dependent variable make a difference to the estimates. In fact whether one uses Davis et al. (2007)-style growth rates, simple percentage changes or log difference is of little consequence for the results.

In the third column I investigate whether the contemporaneous-timing-distribution might be leading to bias if the instrument is capturing local economic development. Specifically, exogeneity requires that variation is driven by people refinancing at particular times for reasons independent of local conditions in the run-up to 2008. In fact, using the two-year lag weighting function produces similar results to the main specification, though comes as the cost of dropping observations because Land Registry data are only available for England and Wales.

It is also worth noting why this predicted cash-flow shock only affects mortgagors and not the owners of the locally non-tradable establishments. First, small businesses predominantly took variable-rate bank loans before the Crisis and only a few larger businesses were able to access fixed-rate capital-market instruments. In 2007, 99.9% of firms in the locally non-tradable sector employed fewer than 1,300 people and accounted for over half of all locally non-tradable employment.[42] Second, more than 60% of company directors live in a different region to where their establishments are located.[43] So it seems unlikely that the results are driven by directors taking the windfall from their own mortgage to support employment.

In the fourth column I examine the contribution of leverage to locally non-tradable employment growth. Two things stand out. First, the large coefficient associated with the loan to income (LTI) ratio suggests that neighborhoods that were highly indebted

[41] For example, if non-conforming, financially vulnerable households are more likely to take out an adjustable-rate mortgage in the US, and it is not fully controlled for, then this could lead to downward bias in the coefficient estimates.

[42] While Gürkaynak et al. (2019) exploit variation in the type of debt businesses took out to answer a similar question, stock-exchange listed businesses are a small fraction of overall employment. Some firms finance their activities with trade-credit, which might have complex relationships with policy rate pass-through. There is no obvious spatial variation in this behavior.

[43] See Bahaj et al. (2017) for details.

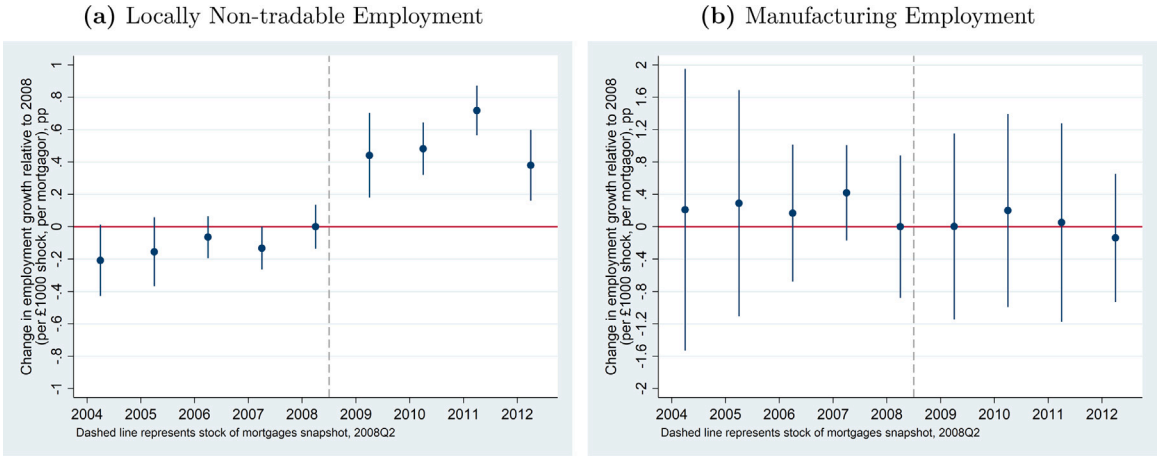

Fig. 11. Event-study Coefficient Estimates. These figures show the main coefficients from IV regressions analogous to the main specification. The dependent variable is the change in the three-year employment growth, up until the date shown, relative to 2008. Since the 2008 coefficient is very close to zero in both figures in the regressions, the coefficient in the third column of Table 5 is very close to the 2010 data point in the left-hand panel. Vertical lines show 90% confidence intervals.

experienced worse employment outcomes during the Great Recession. Comparing the cash-flow and LTI coefficients, and the interquartile ranges in Table 4, suggests that the negative indebtedness shock might have only been partially offset from the positive cash-flow shock in these areas.[44] The second thing that stands out is the interaction term, which is not statistically different from zero. This is potentially surprising as one might expect the marginal impact of a cash-flow boost to be higher for neighborhoods with higher LTIs (as they are likely to have higher marginal propensities to consume).

Another concern is that the neighborhoods that received the largest cash-flow shocks are in some ways different, perhaps because of their employment structure. In the fifth column I add in macro controls for the GVA per capita, the share of manufacturing employment and the overall employment change (excluding locally non-tradable employment) between 2007 and 2010. Most of this variation is picked up by the fixed effects and the estimate for column 5 is close to the main specification.[45]

Another plausible channel through which local non-tradable employment could be supported is via property prices. The first concern is that neighborhoods with higher house prices received larger cash-flow shocks. While this is true to a limited extent, property prices are spatially correlated so much of this level effect is soaked up in the local-authority fixed effect.[46] In addition, I add in the local authority house price index in the sixth column to control for prices at a more granular level. The second concern is that house price falls led to wealth or collateral effects that caused some of the variation in consumption. I therefore add in the change in house prices between 2008Q4 and the end of 2010Q2 in column 6, which leads to a small increase in the central estimate.[47] London had quite unique property price dynamics and also accounts for quite a bit of the variation in cash-flow shock and employment. Inner London wards are likely to be very different compared to the rest of the country. The seventh column builds upon the controls and shows that dropping these does little to the central estimate.[48]

It is also worth probing the timing of the specification I have chosen. Fig. 11 shows an event study of the main IV regression coefficients for alternative three-year windows, relative to 2008, starting from 2001–2004. The left-hand panel shows that the regressions provide similar estimates before 2008.[49] After rising to around 0.7pp in 2011, the estimates fall back to zero after the cash-flow shock had almost fully passed through to households by 2012. The choice of the 2007 to 2010 window is a trade-off. On the one hand it gives establishments time to respond to the heterogeneous consumption patterns. On the other, it is important to capture the pre-crisis status of employment and later estimates are more likely to be compromised by second-round effects of the cash-flow shock (such as the consumption response of those whose employment status had been affected by the initial cash-flow shock) and the broader effects of monetary policy. The right-hand panel shows the analogous estimates for manufacturing employment, which are volatile but close to zero for most of the pre- and post-2008 sample. The manufacturing comparison also helps to rule out other channels such as labor-supply effects or director windfalls.

[44] For more discussion on the effects of monetary policy in the presence of high household indebtedness, see Cumming and Hubert (2021) and Cumming and Hubert (2022).

[45] Fixed effects should also take account of the unobserved fraction of households that experience negative cash-flow effects due to their large liquid balances.

[46] For example, the average house price for a first-time buyer in the first six months of 2008 was £266,000 in London compared to £118,000 in the North of England.

[47] Residential property price changes are likely to be correlated with local commercial property prices too, though most small firms do no own property.

[48] Inner London includes: City, Camden, Greenwich, Hackney, Hammersmith and Fulham, Islington, Kensington and Chelsea, Lambeth, Lewisham, Southwark, Tower Hamlets, Wandsworth and Westminster.

[49] Given that Bank Rate was monotonically increasing from the middle of 2005, the slight negative coefficients are to be expected in the pre-period.

Finally, Table A9 in the Appendix shows some further robustness checks that look at the importance of observation heterogeneity. The first column shows that the central estimates fall by 10bp when the regressions are equally weighted across wards, though dropping the top quartile of wards according to locally non-tradable employment does not change the main coefficient by much. The third column shows the estimates lose some accuracy when dropping the wards with the very largest predicted cash-flow shocks. As discussed in Verner and Gyongyosi (2018), these results likely partly reflect the fact that the effects are stronger in larger and more densely populated areas, where more of the windfall is spent locally. Larger cash-flow shocks are also more likely to lead to discrete behavioral changes so it is unsurprising that the third column demonstrates that cash-flow shock variation is important.[50] These factors are another justification for dropping the most central London wards as a cross-check in Table 6.[51] Overall, the additional robustness checks suggest that neighborhoods with the largest predicted cash-flow shocks are important for accurate identification, partly because instrument variation is limited by a regression to the mean of mortgaging behavior across neighborhoods. That said, the fact that point estimates remain reasonably stable across specifications suggests that the employment effect is indeed apparent in the data.

5.4. The cash-flow multiplier

The cash-flow-employment estimates in this study are quantitatively precise and work through a direct consumption channel. A useful way to benchmark these results is therefore to consider what they imply for the cash-flow shock required to add, or sustain, one job in the economy.

Each neighborhood has around 500 mortgagors so a £1,000 cash-flow shock leads to an aggregate income boost of around £500,000. Each neighborhood employs around 500 people in the locally non-tradable sector, which means my estimates suggest this cash-flow shock increased employment by around 4 people-years over the period of study. Each locally non-tradable job therefore required a total cash-flow boost of around £125,000 (or $180,000 in 2008 PPP). There is a large body of evidence that provides plausible ranges for the typical marginal propensities to consume for these mortgagors. Survey evidence from the US (Fuster et al., 2021), the UK (Bunn et al., 2018) and the Netherlands (Christelis et al., 2019) points to a typical marginal propensity to consume out of a positive income shock for mortgagors of around 0.2.[52] Finally, between 2006 and 2017 the average mortgagor spent around half their total consumption basket on locally non-tradable goods and services.[53] The implied *direct* spend of £12,500 for each locally non-tradable job saved is reassuringly close to the average locally non-tradable wage of around £15,000 per year.[54]

The headline cost-per-job of $180,000 I find is around six times larger than that found in Verner and Gyongyosi (2018), though this is likely an upper bound given the difference in wages between Hungary and the UK. To the extent there is a like-for-like difference, it likely reflects their focus on aggregate unemployment, which captures some general equilibrium effects such as the role of collateral and spillovers to other industries not narrowly defined here as locally non-tradable.[55] Moreover, there is a range of evidence suggesting we might expect a *negative* shock to have a larger effect on households if it takes them closer to their borrowing constraints (e.g., Christelis et al. (2019)). In a similar vein, the negative wealth shock explored in Mian and Sufi (2014) leads to an economically large effect. They estimate that a one standard deviation (8.5%) net wealth shock reduced locally non-tradable employment by around 3.1pp, compared to this paper, which finds that a one standard deviation £1,260 cash-flow shock increased locally non-tradable employment growth by 0.55pp.[56]

A key contribution of this work is therefore to contextualize the findings of the mortgage design literature. Various studies have called for mortgage market reform along the lines of the UK's comprehensive adjustable-rate market (e.g., Campbell et al. (2018); Greenwald et al. (2018); and Guren et al. (2018)). This paper shows that the UK indeed benefited from a mortgage market that provided large-scale cash-flow relief during the Great Recession. But the comparison to Mian and Sufi (2014) and Verner and Gyongyosi (2018) suggests that aggregate indebtedness and volatility of net worth might be quantitatively much more important for employment volatility.[57] Macroprudential policy and mortgage reform that focuses on minimizing balance sheet (rather than cash-flow) volatility might therefore achieve better overall employment stabilization in future recessions.

[50] For example, households may think of discretionary purchases as a binary choice rather than adjusting them on the margin when the alternative is saving.

[51] See also Fig. 7.

[52] An older literature using cross-sectional survey responses following fiscal stimulus finds similar results for non-durable goods and highlights the importance of liquidity constraints (e.g., Johnson et al. (2006); Agarwal et al. (2007); Parker et al. (2013)). There is currently little evidence regarding how consumption decisions are split across consumption categories, or how much people might spend on locally-provided goods and services. In Appendix 1.7 I use UK expenditure surveys and tax rebate shocks to provide some supporting evidence for how windfalls are spent on different parts of the consumption basket.

[53] Data from the Living Costs and Food Survey. Locally non-tradable consumption is defined as spending on restaurants, recreation, food, and miscellaneous goods and services.

[54] Importantly, £12,500 is between the average wage of food-sector workers (£10,700) and retail-sector workers (£13,400). See Table A7 in the Appendix.

[55] More generally, the fiscal multipliers surveyed by Chodorow-Reich (2019) are difficult to compare to the partial-equilibrium effects in this paper. But the preferred estimate of $50,000 per job in that work is consistent with the direct effect I find from local consumption assuming a marginal propensity to consume of around 0.4, which is in line with cash-flow-consumption estimates of papers such as Agarwal et al. (2019). General-equilibrium effects and the *missing intercept* problem of Wolf (2019) makes it difficult to draw out the conclusions for aggregate monetary policy. Nevertheless, the employment cash-flow channel explored here does support recent evidence on the monetary transmission mechanism (e.g., Holm et al. (2020); Almgren et al. (2019); and Cumming and Hubert (2021) and theoretical work using heterogeneous models (e.g., Kaplan et al. (2018)).

[56] Again, ignoring the *missing intercept* problem, the mean net wealth shock of Mian and Sufi (2014) implies a locally non-tradable response of 2.4pp vs 1.3pp in this study. Meaningful comparisons to Mian and Sufi (2014) of course rely on assuming the distribution of shocks is approximately similar in both cases.

[57] Indeed, more recent evidence suggests that household balance sheets might be central to large swings in consumption and employment (e.g., Kaplan et al. (2017); Kehoe et al. (2019) and Garriga and Hedlund (2020)).

Table 7
Neighborhood regressions by establishment type.

	Employment growth, pp								
	Industrial sector				Establishment employment				
	Food	Shops	Cars	Single-unit	1-5	6-13	14-36	37-162	163+
Cash-flow shock (£000)	0.925*** (7.701)	0.080 (0.746)	0.305 (0.544)	1.005*** (4.417)	0.490*** (6.550)	0.201* (2.105)	0.662*** (8.102)	0.623*** (7.430)	0.060 (0.169)
Local-authority fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Specification	IV	IV	IV	IV	IV	IV	IV	IV	IV
Observations	9,045	9,062	8,606	9,068	9,067	8,976	8,100	5,059	2,047
R^2	0.014	0.000	0.000	0.016	0.006	0.001	0.004	0.002	0.000
	Characteristics of economy-wide establishment buckets								
Number of establishments	155,339	280,970	48,475	346,728	314,565	109,626	44347	13,839	2,407
Number of employees	1,306,245	2,709,070	406,516	1,604,763	801,433	924,283	926,373	887,197	882,545

Note: *p<0.05; **p<0.01; ***p<0.001.
The dependent variable is the neighborhood-level employment growth between 2007 and 2010 for establishments matching the column categories in 2007. Establishment size buckets approximately match employment quintiles. Regressions are weighted by employment and standard errors are clustered at the region level (*t*-statistics shown in parentheses). There are 10 regions and 389 local authorities. The economy-wide characteristics are from 2007.

6. Decomposition of employment responses

What does it mean for cash-flow shocks to generate employment growth in the locally non-tradable sector? This section provides more evidence about which sectors, establishment sizes and margins of adjustment led to the cross-sectional differential effects discussed above.

6.1. Sectors and sizes

The basis for looking at the employment response of locally non-tradable establishments rests on the idea that these businesses rely on local consumption to generate their revenues and, ultimately, support employment. The definition used for locally non-tradable businesses in this paper can be broken down further into sic-codes based on food, shops and cars (sales and service). In this subsection I exploit the establishment-level data to run the main specification of Eq. (3). I use the cash-flow instrument on neighborhood employment for subsets of locally non-tradable businesses to investigate exactly how cash-flow shocks affect employment behavior.

The first three columns of Table 7 show the employment growth responses when restricting establishment sic-codes to the three distinct sectors comprising locally non-tradable businesses. As before, the number of observations refers to the number of neighborhoods, so the last two rows of the table show the 2007 economy-wide statistics relevant for each column. It shows that locally non-tradable employment accounted for around 4.4 m employees in the UK, spread across around 500,000 establishments. The coefficients reported in the first three columns suggest that locally-non tradable employment growth is almost entirely driven by the food sector.[58] To the best of my knowledge, this fact has not been pointed out elsewhere.[59] But the result is also not particularly surprising given that the food sector has been the focus of numerous studies examining the flexibility of its labor market (e.g., Card and Krueger (2000); Aaronson and French (2007) and Dube et al. (2016)). Moreover, it is plausible that a restaurant that is twice as busy needs more waiting staff. It is less clear that is true for check-out staff at a shop.[60] Since many food outlets are franchises, it is also not surprising that the fourth column shows a significant response of single-establishment businesses (i.e. where there is no establishment network).

The final five columns of Table 7 split locally non-tradable employment responses by establishment size in 2007. The buckets each account for roughly a fifth of locally non-tradable employment so the largest category only contains around 2,400 of the largest establishments.[61] The main impression from the coefficients is that much of the overall employment response comes from medium-sized establishments. Given the sectoral heterogeneity, however, the rest of this section dives deeper into the food sector to understand the employment drivers of cash-flow shocks.

To start, Table 8 shows the results of the usual neighborhood regressions when including only food-sector employment. The final rows show that just over 80% of all food-sector establishments constitute the entire business and those units employ just over half of food-sector workers. Of more direct interest, the coefficient in the first column suggests that employment responses for single-unit firms are reasonably close to the overall effect (i.e. compared to the first column of Table 7). The final five columns once again show employment quintiles and the bucket sizes reflect the establishment size distribution of the food sector in 2007. On the face

[58] The food sector includes restaurants, pubs and catering companies. For full details of sic-code definitions see Appendix Table A3.

[59] For example, Mian and Sufi (2014), Di Maggio et al. (2017) and Verner and Gyongyosi (2018) test for a combination of *restaurants and retail* in their robustness exercises. These tests likely masked the fact that the restaurants accounted for much of the variation.

[60] Even if people switch from eating in to out, this might not translate into falls in grocery employment.

[61] Only around a quarter of neighborhoods have these very large establishments, which explains the lower number of observations.

Table 8
Neighborhood regressions for the food sector.

	Employment growth, pp					
		Establishment employment				
	Single-unit	1-5	6–11	12–20	21–44	45+
Cash-flow shock	1.068**	0.280	−0.138	1.522***	0.248*	1.895***
(£000)	(2.758)	(1.343)	(−0.521)	(5.707)	(2.374)	(5.736)
Local-authority fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Specification	IV	IV	IV	IV	IV	IV
Observations	8,209	8,175	7,836	6,490	4,822	2,262
R^2	0.011	0.001	0.000	0.008	0.000	0.012
	Characteristics of economy-wide establishment buckets					
Number of establishments	121,402	91,630	32,100	17,171	9,270	2,980
Number of employees	710,822	256,136	250,672	256,512	269,572	264,435

Note: *p<0.05; **p<0.01; ***p<0.001.
The dependent variable is the neighborhood-level employment growth between 2007 and 2010 for food-sector establishments. Establishment size buckets approximately match employment quintiles. Regressions are weighted by employment and standard errors are clustered at the region level (*t*-statistics shown in parentheses). There are 10 regions and 389 local authorities. The economy-wide characteristics are from 2007.

Table 9
Extensive margin of food-sector employment.

	Establishment Birth or Death Rate, pp				
	Establishment employment				
	1–5	6–11	12–20	21–44	45+
Death Rates					
Cash-flow shock	−0.134	−0.312	−0.812***	−0.006	0.073
(£000)	(−1.439)	(−1.951)	(−9.177)	(−0.056)	(0.558)
Local-authority fixed effects	Yes	Yes	Yes	Yes	Yes
Specification	IV	IV	IV	IV	IV
Observations	91,268	29,902	15,687	9,595	2,786
R^2	0.000	0.001	0.006	0.000	0.000
	Death rate of establishment buckets, 2007–2010				
	0.520	0.435	0.405	0.274	0.267
Birth Rates					
Cash-flow shock	0.007	0.024	0.052	0.202**	0.500***
(£000)	(0.148)	(0.208)	(0.792)	(2.736)	(6.954)
Local-authority fixed effects	Yes	Yes	Yes	Yes	Yes
Specification	IV	IV	IV	IV	IV
Observations	113,614	34,487	29,713	16,673	7,251
R^2	0.000	0.000	0.000	0.001	0.006
	Birth rate of establishment buckets, 2007–2010				
	0.283	0.227	0.196	0.172	0.147

Note: *p<0.05; **p<0.01; ***p<0.001.
The dependent variable is the change in 2007–2010 establishment birth or death rate, defined relative to comparative food-sector establishments in 2007. Standard errors are clustered at the region level (*t*-statistics shown in parentheses). There are 10 regions and 389 local authorities.

of it, the coefficients reveal that the overall cash-flow response is driven by two buckets: the mid-size establishments employing 12 to 20 people, and the largest establishments employing more than 45. For the latter group, the estimates imply that an average neighborhood cash-flow boost of £1,000 per mortgagor led to an increase in employment growth of almost 2pp.

6.2. The extensive margin

The fact that the vast majority of regressions in this paper consider the growth rate of employment as the dependent variable makes it tempting to frame the problem in terms of a decision to hire or fire an additional worker. But overall employment growth can equally be driven by a reduction in the establishment death rate or an increase in the establishment birth rate.[62] Whether employment effects are driven by the intensive or extensive margin is ultimately an empirical question. Neighborhood-level regressions can often mask which margin of adjustment is most important.

[62] Business demographics are well documented and share many similarities across countries (e.g., Bartelsman et al. (2005)).

Table 10
Intensive margin of food-sector employment.

	Employment growth, pp				
	Establishment employment				
	1–5	6–11	12–20	21–44	45+
Birth Rates					
Cash-flow shock	1.975***	1.204***	0.730***	0.068	−0.224
(£000)	(13.258)	(9.434)	(8.261)	(0.382)	(−0.610)
Local-authority fixed effects	Yes	Yes	Yes	Yes	Yes
Specification	IV	IV	IV	IV	IV
Observations	8,171	6,130	4,511	3,123	1,046
R^2	0.044	0.028	0.012	0.000	0.000
	Characteristics of economy-wide establishment buckets				
Number of establishments	43,381	17,891	10,809	6,708	2,189
Number of employees	125,601	140,764	163,520	195,602	176,308

Note: *p<0.05; **p<0.01; ***p<0.001.
The dependent variable is the neighborhood-level employment growth between 2007 and 2010 for food-sector establishments that traded in both years. Regressions are weighted by employment and standard errors are clustered at the region level (*t*-statistics shown in parentheses). There are 10 regions and 389 local authorities. The economy-wide characteristics are from 2007.

To investigate the extensive margin I run linear probability models at the establishment level using the neighborhood-level instrument developed above, once again grouping establishments into the five size buckets.[63] The coefficients reported from these can be interpreted as percentage point changes in the birth or death rate following a £1,000 average cash-flow boost, as in the main specifications. Assuming the marginal effect is homogeneous within the bucket allows one to also interpret the coefficient as the contribution to the total employment effect shown in Table 8, where the sign is reversed for the death rate.[64] Regression coefficients are reported in Table 9 and plotted in Fig. 12 for ease of comparison.

The top panel of Table 9 shows that the three-year death rate for establishments in the food sector ranges from around 52% for the smallest to just over 25% for the largest. The fact that failure rates are often declining in establishment size is well documented (e.g., Bahaj et al. (2019)), though many earlier studies have focused on the manufacturing sector (e.g., Dunne et al. (1989)). But food-sector death rates are some of the highest across industries, which means the sensitivity of the extensive margin is especially important for overall employment growth.

There has been little research looking at the responsiveness of death rates to precisely calibrated demand shocks.[65] The establishment exit regressions employed here include all establishments that existed in 2007 and define the dependent variable as a dummy variable equal to 1 if the establishment was recorded as dying at some point in the next three years. Interestingly, the top panel of Table 9 suggests there is a substantial impact on the death rate of medium-sized food-sector establishments, but only for that employment quintile. Taken at face value, around half of the overall employment effect for establishments employing 12 to 20 people comes from a reduction in the establishment death rate.

The lower panel of Table 9 examines establishment birth rates, defined as the number of new establishments formed between 2007 and 2010 as a fraction of the number that existed in 2007. The final row shows that the three-year establishment birth rate also falls by half when moving from the smallest to the largest establishments, though is lower than the death rate across all quintiles. The cash-flow effect coefficients show the interesting pattern that birth rates are most affected for the largest establishments. The 0.5pp estimate in the last column is also likely to understate the contribution to overall employment growth if the effect is increasing in establishment size.[66] The lesson from this panel is that cash-flow shocks incentivize existing firms to open up new outlets rather than encouraging new entrepreneurs, as very few businesses launch with more than 45 employees.[67]

6.3. The intensive margin

For the establishments that survived the recession, it is worth investigating to what extent cash-flow shocks contributed to their employment growth on the intensive margin. Table 10 shows the results of the neighborhood-level regressions when restricting the

[63] Linear probability models make it easier to cluster standard errors at the regional level but instrumental-variables probit models deliver very similar marginal effects. Results available upon request.

[64] For example, a coefficient of −1 for the death rate means that 1pp fewer establishments died per £1,000 cash-flow shock between 2007 and 2010. That is roughly equivalent to employment growth over the same time period increasing by 1pp.

[65] There is some evidence of a correlation between house prices and unemployment and the extensive margin (e.g., Moscarini and Postel-Vinay (2012) and Fort et al. (2013)). There is also an extensive literature looking at product switching following demand shocks (e.g., Antoniades et al. (2018)) or restaurant exit in the face of changes to wages (e.g., Luca and Luca (2019)). Aaronson et al. (2018) use a putty-clay framework to show that minimum wage increases led to increased entry and exit in the food sector. In that model, establishments are free to choose their input mix before opening up but must follow a Leontief production function once in place.

[66] For example, imagine an extreme example where a £1,000 cash-flow shock leads to only very large establishments opening up. In that case the establishment birth rate, which is relative to the bucket as a whole, understates the new employment generated.

[67] This can be seen as a demand-shock analogy to the finding that minimum wage laws deter new establishments from locating in an neighborhood (e.g., Rohlin (2011)). This result is also consistent with the result that decentralized firms tend to do better following macro shocks (e.g., Aghion et al. (2017)).

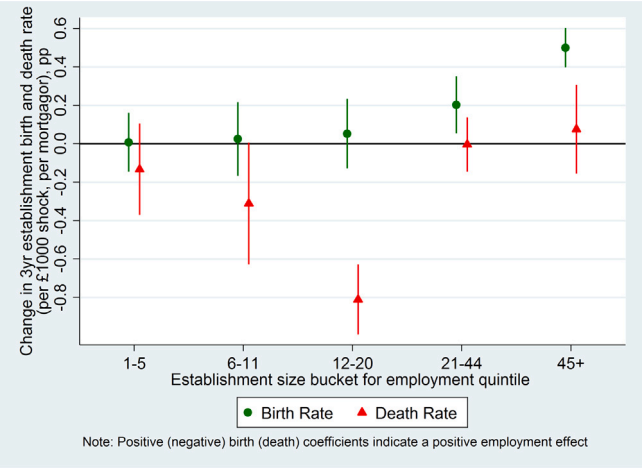

Fig. 12. The Drivers of Food-sector Employment. This figure plots the marginal extensive-margin effects between 2007 and 2010 using the cash-flow instrument and coefficient data from Fig. 12. The five employment buckets each account for around a fifth of food-sector employment. Vertical lines show 90% confidence intervals.

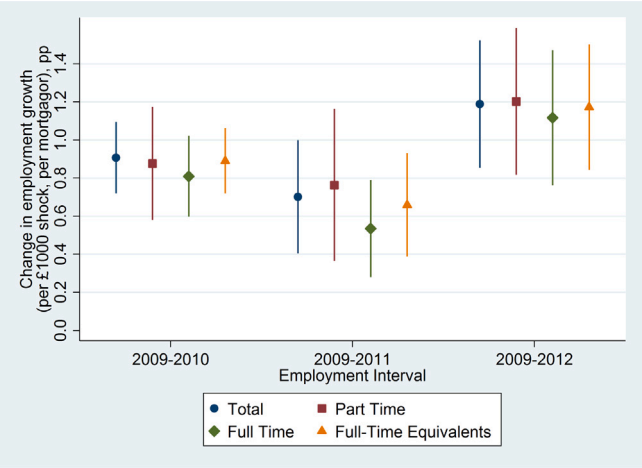

Fig. 13. Full-time and Part-time Employment. This figure plots the marginal cash-flow employment effect using the cash-flow instrument on a consistent basis with the main specification. The four estimates substitute the main dependent variable with different types of employment captured in the BRES survey (around 50,000 food-sector establishments employing 750,000). The BRES survey starts in 2009 and so the time period cannot fully replicate the central specification, so three alternatives are used. Vertical lines show 90% confidence intervals.

sample to food-sector establishments that were trading in both 2007 and 2010. It shows that the cash-flow effect on marginal employment decisions was largest for the smallest establishments. For example, an average £1,000 cash-flow shock increased employment growth for the very smallest establishments by almost 2pp. For food-sector establishments employing more than 20 people, the estimated effects are small and statistically insignificant.

Of course, the number of hired workers is not the only margin of adjustment available to businesses.[68] Using a survey that launched in 2009 it is possible to look at how the composition of part-time and full-time workers changed in response to cash-flow shocks.[69] One hypothesis to test is whether firms adjusted the number of hours as well as the number of heads. Fig. 13 plots the coefficients of the specification where the dependent variable includes all employees, part-time, full-time and full-time-equivalent workers in the food sector. Given the survey only starts in 2009 it is not possible to run the regressions over the usual 2007–2010 window, so I show three alternatives. None of the three specifications provide much evidence for substitution between full-time and part-time workers. In fact, employment effects are very similar for all types of workers, suggesting that establishments do not

[68] I found no evidence of changes in wages associated with cash-flow shocks, which is likely driven by the fact that the food-sector pays wages that are often very close to the minimum wage. This has also been documented elsewhere (e.g., Giroud and Mueller (2019)).

[69] In September 2009 the ONS started the Business Register and Employment Survey (BRES) to improve the accuracy of establishment and sectoral employment data.

respond to these cash-flow shocks by altering the composition of their staff.[70] Nevertheless, one limitation of the survey is that, although it captures a large share of employment, it only captures larger establishments and so might hide composition effects for the smaller outlets.

6.4. Discussion

So what drove the cash-flow-effect employment growth in food establishments? I find that the overall effect of 0.9pp can be decomposed into roughly three equal effects. First, around a third of the overall effect was driven by the smallest establishments expanding the number of workers they employed. The small and insignificant coefficients in the first two columns of Table 8 mask the fact that many establishments moved between employment buckets due to this intensive-margin employment growth. Specifically, the intensive-margin coefficient on establishments employing fewer than 11 employees is around 1.6pp. Approximately half of establishments this size survived the period and jobs in these establishments made up 40% of food-sector employment. That means the intensive margin for these establishments accounts for around 0.3pp of the overall employment growth between 2007 and 2010.

Second, employment growth from mid-sized establishments that employed between 12–20 employees in 2007 made up another 0.3pp. Around half of that was on the intensive margin (i.e. the 0.73 coefficient in Table 10, for a fifth of employment) and the other half was on the extensive margin via a counter-factually lower probability of establishment death. Put another way, establishments in this bracket were especially vulnerable to closure and positive cash-flow shocks in the surrounding area resulted in turnover boosts that marginally increased their probability of survival.

Third, positive neighborhood cash-flow shocks increased the probability of very large establishments launching. The negative intensive-margin effect, the tiny death-rate effect and the downward bias of the birth-rate coefficient suggests that the birth-rate effect actually makes up most of the 1.9pp effect for largest establishments in Table 8. Once again representing a fifth of employment, that suggests this effect accounts for a little over a third of the overall employment response of the food-sector.

7. Conclusion

I find that the monetary easing by the Bank of England in the fall of 2008 had a significant and immediate cash-flow impact on people with an adjustable-rate mortgages. In areas where the overall cash-flow shock was especially large, it supported consumption, including consumption for locally provided goods and services, thereby supporting employment in these sectors. Although monetary policy works through multiple channels, I find that the cash-flow channel offset some important adverse shocks during the Great Recession.

My finding that a cash injection to the economy of £125,000 directly saves one locally non-tradable job reminds us that monetary policy is powerful, though this metric likely underestimates the overall support provided. The large employment response I find in the restaurant sector reflects the flexibility of its labor market, but other locally non-tradable firms surely benefit from increased spending. Wider spill-overs to the tradable sector are also likely important even if they are more difficult to measure. In addition, the cost-per-job metric hides the fact that the total cash-flow boost imparted by monetary policy is so large because of the sheer scale of the mortgage market.

One implication of this work is that monetary policy can lead to heterogeneous employment effects across the country. Neighborhoods that happened to contain a large proportion of people on adjustable-rate mortgages and employed a large fraction of their labor force in the food sector benefited the most from the first-round effects of accommodative monetary policy. But my work also has implications for the traction of monetary policy over time. Since policy rates reached their Great-Recession nadir in 2009, many countries have seen an increase in the demand for fixed-rate contracts, as people have attempted to lock in low mortgage rates for an extended period (e.g., Cumming (2019)). This could change the monetary transmission mechanism in important ways going forward. To the extent that many people are shielded from interest rate rises in the near future and the cash-flow channel is a key part of transmission, we would not expect to see symmetric potency of monetary policy in a tightening cycle, all else equal.

Finally, although this work is motivated by understanding the cash-flow channel of monetary policy, it has implications for the transmission of other shocks that also lead to consumption changes. It so happens that the food sector (and locally non-tradable sector more broadly) contains some of the businesses most exposed to precautionary reductions in demand and mandatory social-distancing measures stemming from the Coronavirus public health crisis. This paper provides some indication of how such consumption shocks might translate into employment outcomes for different regions, sectors and establishment sizes.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.euroecorev.2021.104019.

[70] Horton (2017) shows that the intensive margin of adjustment is small following the introduction of minimum wages though there is some substitution of low for high productivity workers. A comparable investigation is beyond the scope of this paper. See Footnote 6.2 for a Leontief production function analogy.

References

Aaronson, Daniel, French, Eric, 2007. Product market evidence on the employment effects of the minimum wage. J. Labor Econ. 25 (1), 167–200.

Aaronson, Daniel, French, Eric, Sorkin, Isaac, To, Ted, 2018. Industry dynamics and the minimum wage: A putty-clay approach. Internat. Econom. Rev. 59 (1), 51–84.

Agarwal, Sumit, Chomsisengphet, Souphala, Liu, Chunlin, Souleles, Nicholas S, 2015. Do consumers choose the right credit contracts? Rev. Corp. Finance Stud. 4 (2), 239–257.

Agarwal, Sumit, Deng, Yongheng, Gu, Quanlin, He, Jia, Qian, Wenlan, Ren, Yuan, 2019. Mortgage debt, hand-to-mouth households, and monetary policy transmission. Hand-To-Mouth Households, And Monetary Policy Transmission (January 29, 2019).

Agarwal, Sumit, Liu, Chunlin, Souleles, Nicholas S, 2007. The reaction of consumer spending and debt to tax RebatesEevidence from consumer credit data. J. Political Econ. 115 (6), 986–1019.

Agarwal, Sumit, Mazumder, Bhashkar, 2013. Cognitive abilities and household financial decision making. Am. Econ. J. Appl. Econ. 5 (1), 193–207.

Aghion, Philippe, Bloom, Nicholas, Lucking, Brian, Sadun, Raffaella, Van Reenen, John, 2017. Turbulence, Firm Decentralization and Growth in Bad Times. Technical Report, National Bureau of Economic Research.

Aikman, David, Bridges, Jonathan, Kashyap, Anil, Siegert, Caspar, 2019. Would macroprudential regulation have prevented the last crisis? J. Econ. Perspect. 33 (1), 107–130.

Almgren, Mattias, Gallegos, Jose E, Kramer, John, Lima, Ricardo, 2019. Monetary policy and liquidity constraints: Evidence from the Euro area. Available at SSRN 3422687.

Antoniades, Alexis, Clerides, Sofronis, et al., 2018. Micro-Responses to Shocks: Pricing, Promotion, and Entry. Technical Report, CEPR Discussion Papers.

Aron, Janine, Muellbauer, John, 2016. Modelling and forecasting mortgage delinquency and foreclosure in the UK. J. Urban Econ. 94, 32–53.

Badarinza, Cristian, Campbell, John Y, Ramadorai, Tarun, 2017. What calls to arms? International evidence on interest rates and the choice of adjustable-rate mortgages. Manage. Sci..

Bahaj, Saleem A, Foulis, Angus, Pinter, Gabor, 2017. Home Values and Firm Behaviour. Technical Report, Centre for Macroeconomics.

Bahaj, Saleem, Pinter, Gabor, Foulis, Angus, Surico, Paolo, 2019. Employment and the Collateral Channel of Monetary Policy. Bank of England Working Paper.

Bartelsman, Eric, Scarpetta, Stefano, Schivardi, Fabiano, 2005. Comparative analysis of firm demographics and survival: evidence from micro-level sources in OECD countries. Ind. Corp. Chang. 14 (3), 365–391.

Beraja, Martin, Fuster, Andreas, Hurst, Erik, Vavra, Joseph, 2018. Regional heterogeneity and the refinancing channel of monetary policy. Q. J. Econ. 134 (1), 109–183.

Bernanke, Ben S., 2007. The Financial Accelerator and the Credit Channel. Speech 296. Board of Governors of the Federal Reserve System.

Best, Michael Carlos, Cloyne, James, Ilzetzki, Ethan, Kleven, Henrik Jacobsen, 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.

Best, Michael Carlos, Cloyne, James, Ilzetzki, Ethan, Kleven, Henrik, 2018. Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches. Technical Report, National Bureau of Economic Research.

Bunn, Philip, Rostom, May, Domit, Silvia, Worrow, Nicola, Piscitelli, Laura, 2013. The financial position of british households: Evidence from the 2013 NMG consulting survey. Bank Engl. Q. Bull. 53 (4), 351–360.

Bunn, Philip, Roux, Jeanne Le, Reinold, Kate, Surico, Paolo, 2018. The consumption response to positive and negative income shocks. J. Monet. Econ. 96, 1–15.

Byrne, David, Kelly, Robert, O'Toole, Conor, 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, Alessandro, Monacelli, Tommaso, Stracca, Livio, 2013. Housing finance and monetary policy. J. Eur. Econ. Assoc. 11 (suppl_1), 101–122.

Campbell, John Y, Clara, Nuno, Cocco, Joao F, 2018. Structuring mortgages for macroeconomic stability. Available at SSRN 3135366.

Campbell, John Y, Cocco, Joao F, 2003. Household risk management and optimal mortgage choice. Q. J. Econ. 118 (4), 1449–1494.

Card, David, Krueger, Alan B, 2000. Minimum wages and employment: a case study of the fast-food industry in new jersey and pennsylvania: reply. Amer. Econ. Rev. 90 (5), 1397–1420.

Cava, Gianni La, Hughson, Helen, Kaplan, Greg, 2016. The Household Cash Flow Channel of Monetary Policy. RBA Research Discussion Papers rdp2016-12, Reserve Bank of Australia.

Chatterjee, Satyajit, Eyigungor, Burcu, 2019. The Firm Size and Leverage Relationship and Its Implications for Entry and Concentration in a Low Interest Rate World. FRB of Philadelphia Working Paper.

Chodorow-Reich, Gabriel, 2019. Geographic cross-sectional fiscal spending multipliers: What have we learned? Am. Econ. J. Econ. Policy 11 (2), 1–34.

Christelis, Dimitris, Georgarakos, Dimitris, Jappelli, Tullio, Pistaferri, Luigi, Van Rooij, Maarten, 2019. Asymmetric consumption effects of transitory income shocks. Econ. J. 129 (622), 2322–2341.

Cloyne, James, Ferreira, Clodomiro, Surico, Paolo, 2020. Monetary policy when households have debt: new evidence on the transmission mechanism. Rev. Econ. Stud. 87 (1), 102–129.

Cloyne, James, Huber, Kilian, Ilzetzki, Ethan, Kleven, Henrik, 2019. The effect of house prices on household borrowing: a new approach. Amer. Econ. Rev. 109 (6), 2104–2136.

Crouzet, Nicolas, Mehrotra, Neil, 2017. Small and Large Firms Over the Business Cycle. Working Paper.

Cumming, Fergus, 2019. Bitesize: Fixing ideas - the slowing of interest-rate pass-through to mortgagors. Bank Underground Blog.

Cumming, Fergus, Dettling, Lisa J, 2020. Monetary Policy and Birth Rates: the Effect of Mortgage Rate Pass-Through on Fertility. FEDS Working Paper.

Cumming, Fergus, Hubert, Paul, 2021. The distribution of households' indebtedness and the transmission of monetary policy. The Review of Economics and Statistics 1–36.

Cumming, Fergus, Hubert, Paul, 2022. House prices, the distribution of dousehold debt and the refinancing channel of monetary policy. Economics Letters 212, 110280.

Davis, Steven J., Haltiwanger, John, Jarmin, Ron, Miranda, Javier, 2007. Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. In: NBER Macroeconomics Annual 2006. Vol. 21. In: NBER Chapters, National Bureau of Economic Research, Inc, pp. 107–180.

Department for Transport, 2016. National Travel Survey: England 2015. Technical Report, Statistical Release.

Di Maggio, Marco, Kermani, Amir, Keys, Benjamin J., Piskorski, Tomasz, Ramcharan, Rodney, Seru, Amit, Yao, Vincent, 2017. Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. Amer. Econ. Rev. 107 (11), 3550–3588.

Dube, Arindrajit, Lester, T. William, Reich, Michael, 2016. Minimum wage shocks, employment flows, and labor market frictions. J. Labor Econ. 34 (3), 663–704.

Dunne, Timothy, Roberts, Mark J., Samuelson, Larry, 1989. The growth and failure of U. S. Manufacturing plants. Q. J. Econ. 104 (4), 671–698.

Dupor, Bill, Guerrero, Rodrigo, 2017. Local and aggregate fiscal policy multipliers. J. Monet. Econ. 92, 16–30.

Fagereng, Andreas, Holm, Martin Blomhoff, Natvik, Gisle James James, 2019. MPC heterogeneity and household balance sheets. Available at SSRN 3399027.

Flodén, Martin, Kilström, Matilda, Sigurdsson, Jósef, Vestman, Roine, 2020. The Economic Journal 131 (636), 1742–1771.

Fort, Teresa C, Haltiwanger, John, Jarmin, Ron S, Miranda, Javier, 2013. How firms respond to business cycles: The role of firm age and firm size. IMF Econ. Rev. 61 (3), 520–559.

Fuster, Andreas, Kaplan, Greg, Zafar, Basit, 2021. What would you do with $500? Spending responses to gains, losses, news and loans. The Review of Economic Studies 88 (4), 1760–1795.
Fuster, Andreas, Willen, Paul S., 2017. Payment size, negative equity, and mortgage default. Am. Econ. J. Econ. Policy 9 (4), 167–191.
Garriga, Carlos, Hedlund, Aaron, 2020. Mortgage debt, consumption, and illiquid housing markets in the great recession. Amer. Econ. Rev. 110 (6), 1603–1634.
Garriga, Carlos, Kydland, Finn E., Sustek, Roman, 2017. Mortgages and monetary policy. Rev. Financ. Stud. 30 (10), 3337–3375.
Giroud, Xavier, Mueller, Holger M, 2017. Firm leverage, consumer demand, and employment losses during the great recession. Q. J. Econ. 132 (1), 271–316.
Giroud, Xavier, Mueller, Holger M, 2019. Firms' internal networks and local economic shocks. Amer. Econ. Rev. 109 (10), 3617–3649.
Greenwald, Daniel, 2018. The mortgage credit channel of macroeconomic transmission. MIT Sloan Research Paper.
Greenwald, Daniel, Landvoigt, Tim, Van Nieuwerburgh, Stijn, 2018. Financial Fragility with SAM? MIT Sloan Research Paper.
Guren, Adam M, Krishnamurthy, Arvind, McQuade, Timothy J, 2018. Mortgage Design in an Equilibrium Model of the Housing Market. National Bureau of Economic Research.
Gürkaynak, Refet S., Karasoy-Can, Hatice Gökce, Lee, Sang Seok, 2019. Stock Market's Assessment of Monetary Policy Transmission: The Cash Flow Effect. Technical report 628, CFS Working Paper.
Haltiwanger, John, 2012. Job creation and firm dynamics in the United States. Innov. Policy Econ. 12 (1), 17–38.
Haltiwanger, John, Jarmin, Ron S, Miranda, Javier, 2013. Who creates jobs? Small versus large versus young. Rev. Econ. Stat. 95 (2), 347–361.
Holm, Martin Blomhoff, Paul, Pascal, Tischbirek, Andreas, 2020. The Transmission of Monetary Policy under the Microscope. Federal Reserve Bank of San Francisco.
Horton, John J, 2017. Price floors and employer preferences: Evidence from a minimum wage experiment. Available at SSRN 2898827.
Johnson, David S, Parker, Jonathan A, Souleles, Nicholas S, 2006. Household expenditure and the income tax rebates of 2001. Amer. Econ. Rev. 96 (5), 1589–1610.
Kaplan, Greg, Mitman, Kurt, Violante, Giovanni L, 2017. The Housing Boom and Bust: Model Meets Evidence. Working Paper 23694, National Bureau of Economic Research.
Kaplan, Greg, Moll, Benjamin, Violante, Giovanni L, 2018. Monetary policy according to HANK. Amer. Econ. Rev. 108 (3), 697–743.
Kartashova, Katya, Zhou, Xiaoqing, 2019. How does mortgage rate refinancing affect consumer spending and debt repayment? Evidence from Canadian consumers.
Kehoe, Patrick J., Midrigan, Virgiliu, Pastorino, Elena, 2019. Debt constraints and employment. J. Political Econ. 127 (4), 1926–1991.
Luca, Dara Lee, Luca, Michael, 2019. Survival of the Fittest: The Impact of the Minimum Wage on Firm Exit. Working Paper 25806, National Bureau of Economic Research.
Luck, Stephan, Zimmermann, Tom, 2020. Employment effects of unconventional monetary policy: Evidence from QE. J. Financ. Econ. 135 (3), 678–703.
Messer, Todd, Siemer, Michael, Gourio, Francois, et al., 2016. A missing generation of firms? Aggregate effects of the decline in new business formation. In: 2016 Meeting Papers. No. 752. Society for Economic Dynamics.
Mian, Atif, Sufi, Amir, 2014. What explains the 2007 2009 drop in employment? Econometrica 82 (6), 2197–2223.
Moscarini, Giuseppe, Postel-Vinay, Fabien, 2012. The contribution of large and small employers to job creation in times of high and low unemployment. Amer. Econ. Rev. 102 (6), 2509–2539.
Nakamura, Emi, Steinsson, Jon, 2014. Fiscal stimulus in a monetary union: Evidence from US regions. Amer. Econ. Rev. 104 (3), 753–792.
Parker, Jonathan A, Souleles, Nicholas S, Johnson, David S, McClelland, Robert, 2013. Consumer spending and the economic stimulus payments of 2008. Amer. Econ. Rev. 103 (6), 2530–2553.
Piskorski, Tomasz, Seru, Amit, 2018. Mortgage Market Design: Lessons from the Great Recession. Technical Report, Brookings papers on economic activity.
Rohlin, Shawn, 2011. State minimum wages and business location: Evidence from a refined border approach. J. Urban Econ. 69 (1), 103–117.
Rubio, Margarita, 2011. Fixed- and variable-rate mortgages, business cycles, and monetary policy. J. Money Credit Bank. 43 (4), 657–688.
Sharpe, Steven A, 1994. Financial market imperfections, firm leverage, and the cyclicality of employment. Am. Econ. Rev. 84 (4), 1060–1074.
Varum, Celeste Amorim, Rocha, Vera Catarina, 2012. The effect of crises on firm exit and the moderating effect of firm size. Econom. Lett. 114 (1), 94–97.
Verner, Emil, Gyongyosi, Gyozo, 2018. Household debt revaluation and the real economy: Evidence from a foreign currency debt crisis.
Violante, Gianluca, Kaplan, Greg, Weidner, Justin, 2014. The wealthy hand-to-mouth. Brook. Pap. Econ. Activity 45 (1 (Spring)), 77–153.
Wolf, C, 2019. The Missing Intercept in Cross-Regional Regressions. Technical Report, Princeton University Working Paper.