



BANK OF ENGLAND

# Staff Working Paper No. 836

## The role of households' borrowing constraints in the transmission of monetary policy

Fergus Cumming and Paul Hubert

January 2020

This is an updated version of the Staff Working Paper originally published on 20 December 2019

Staff Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee.



BANK OF ENGLAND

# Staff Working Paper No. 836

## The role of households' borrowing constraints in the transmission of monetary policy

Fergus Cumming<sup>(1)</sup> and Paul Hubert<sup>(2)</sup>

### Abstract

This paper investigates how the transmission of monetary policy to the real economy depends on the distribution of household debt. Using an original loan-level dataset covering the universe of UK mortgages, we assess the effect of monetary shocks on aggregate consumption by exploiting time variation in a measure of the proportion of households close to their borrowing constraint. We find that monetary policy is most potent when there is a large share of constrained households. In contrast, we find no evidence that the average level of borrowing relative-to-income of the household sector affects the transmission of monetary policy.

**Key words:** Heterogeneity, distributions, mortgage debt, state-dependence.

**JEL classification:** E21, E52, E58.

---

(1) Bank of England. Email: [fergus.cumming@bankofengland.co.uk](mailto:fergus.cumming@bankofengland.co.uk)

(2) Sciences Po — OFCE. Email: [paul.hubert@sciencespo.fr](mailto:paul.hubert@sciencespo.fr)

This paper contains the views of the authors and not necessarily those of the Bank of England or its committees. All errors remain our own. We thank David Aikman, Christophe Blot, Edouard Challe, James Cloyne, Xavier Ragot, Fergal Shortall and Arlene Wong for helpful discussions. This paper also benefited from comments of seminar participants at OFCE, Banque de France, Bank of England, and at the 2019 GDRE, AFSE, IAAE, EEA and MMF conferences. We also thank an anonymous referee for helpful comments.

The Bank's working paper series can be found at [www.bankofengland.co.uk/working-paper/staff-working-papers](http://www.bankofengland.co.uk/working-paper/staff-working-papers)

Bank of England, Threadneedle Street, London, EC2R 8AH

Email [publications@bankofengland.co.uk](mailto:publications@bankofengland.co.uk)

© Bank of England 2020

ISSN 1749-9135 (on-line)

# 1 Introduction

After a decade of near-zero policy interest rates, there still exists considerable debate about the key facets of the monetary transmission mechanism. It is by now well understood that the mortgage market is a key channel of propagation (e.g., [Cloyne et al. \(2019\)](#)) and there is ample micro-data evidence that households close to their borrowing constraints are often the most responsive to shocks (e.g., [Di Maggio et al. \(2017\)](#) and [Georgarakos and Tatsiramos \(2019\)](#)). Progress has also been made in modelling the role of heterogeneity in explaining aggregate dynamics in the context of monetary policy (e.g., [Debortoli and Galí \(2017\)](#) and [Kaplan et al. \(2018\)](#)). But there has been little empirical exploration of how household-borrowing heterogeneity alters the dynamics of the macro-economy. Our paper gets to the heart of how the distribution of household debt affects the transmission of monetary policy.

[Kaplan et al. \(2014\)](#) and [Bilbiie and Ragot \(2017\)](#) show the importance of illiquid-but-wealthy households in the propagation of macroeconomic shocks because of their relatively high marginal propensity to consume. Using individual mortgage data enables us to focus on this set of households. This paper empirically investigates the extent to which the effects of monetary policy depends on the proportion of households close to their borrowing constraints. Since these households are likely to be most directly affected by changes in interest rates, we use time variation in their relative abundance to assess the importance of the distribution of borrowing for the overall traction of monetary policy on aggregate consumption. Intuitively, we find that monetary policy has its greatest effect on the economy when there is a relatively large share of highly-constrained households.

The main empirical challenge is to measure the proportion of these highly-constrained households. Overall household indebtedness has significantly increased over the past decade as interest rates have fallen. But simple loan-to-income (LTI) distribution moments do a poor job of capturing the degree to which households' personal finances are stretched. We exploit proprietary transaction data covering the universe of UK residential mortgages since 2005 to estimate conditional LTIs via micro-data regressions. We can then compare high-LTI mortgagors on a like-for-like basis across time based on a given threshold cut-off. Our novel state-variable identifies the relative prevalence of households close to their borrowing constraints using all the cross-sectional information available. The granularity of the data allows us to control for macro shocks, credit supply and borrower characteristics with unprecedented levels of precision. Our state-variable has a low correlation with standard metrics of balance sheet stretch and different moments of the overall debt distribution.

We then use time variation in our state variable to explore monetary policy dynamics. Using a rich array of UK time-series data and standard identification techniques to form a set of monetary shocks, we proceed by using local projections to investigate the importance of the most indebted households in the transmission of monetary policy to the rest of the economy. Such local projections have been shown to be robust to miss-specification and are a flexible way to generate impulse responses. In particular, we look at the response of consumption to an interaction of monetary policy shocks and our state variable and examine the resulting impulse response functions.

We find that monetary policy is more powerful when a proportion of households have taken on high debt burdens relative to their resources.<sup>1</sup> This effect likely works through at least two mechanisms. First, when households increase their borrowing relative to income, the mechanical effect of monetary policy on disposable income is amplified. Second, households close to their borrowing constraints are likely to exhibit a high marginal propensity to consume. Consistent with recent increased attention on heterogeneity, we find that our results are driven more by the distribution of debt than the overall rise in indebtedness. Replacing our state-variable with a variable capturing average LTI yields no state-dependence effects of monetary policy.

We also find evidence of some asymmetry in the transmission of monetary policy, consistent with [Tenreyro and Thwaites \(2016\)](#) and [Bunn et al. \(2018\)](#): when the share of constrained households is large, restrictive monetary policy has a larger impact than expansionary policy. Further, we find evidence that the potency of monetary policy depends on recent movements in house prices: we show the effect of restrictive monetary policy is greater when house prices have recently fallen compared to when they have recently increased. This is consistent with evidence on the link between consumption, debt and house prices (see [Berger et al. \(2017\)](#)) and the refinancing channel of monetary policy (see [Eichenbaum et al. \(2018\)](#) and [Beraja et al. \(2019\)](#)). Finally, we present some evidence that GDP and industrial production are also differentially affected by monetary policy conditional on the share of constrained households. This is likely to be driven in part by an accounting effect but might also reflect second-round feedback effects stemming from the interaction of consumer spending and firm behaviour.

A potential concern with our approach is that we may be capturing factors other than the state contingency of monetary policy with regard to borrowing constraints. For instance, in recent decades, the UK has undergone significant structural change along several dimensions including the composition and regulation of mortgages. To address this, our specification controls for changes in the structure, price and quantity dynamics of the mortgage market over our sample. More specifically, we separate out the effect of changes to the average duration of mortgages and the share of floating-rate mortgages in issuance. We also control for the benchmark reset interest rate and the overall flow of mortgage originations.

Our paper adds to a rich literature that has established an important link between household balance sheets and the propagation of shocks, both in terms of assets or mortgage debt (e.g., [Attanasio et al. \(2002\)](#); [Iacoviello \(2005\)](#); [Eggertsson and Krugman \(2012\)](#); [Mian and Sufi \(2011\)](#); [Ragot \(2014\)](#); [Misra and Surico \(2014\)](#) and [Carroll et al. \(2017\)](#)) and in terms of income flows (e.g., [Kaplan and Violante \(2014\)](#); [Auclert \(2019\)](#); and [Greenwald \(2018\)](#)). Our paper builds on the recent works of [Cloyne et al. \(2019\)](#) to dig deeper into the overall effects of monetary policy and separate out the various channels that affect household behaviour.

---

<sup>1</sup>Our main result is robust to various specifications of our state-variable, different definitions of borrowing constraints, excluding the Great Recession from the sample, and alternative monetary shocks.

Our work also relates to the literature linking the structure of the mortgage market and the transmission of shocks. Some recent studies have exploited mortgage market heterogeneity to identify shocks at a granular level (e.g., Di Maggio et al. (2017); Flodén et al. (2017); and Cumming (2018)), while others have taken broader lessons for the role of monetary policy propagation (e.g., Calza et al. (2013); Finck et al. (2018); and Piskorski and Seru (2018)). In this study we therefore exploit the unique structure and data granularity of this database of the UK mortgage market as a way to understand the channels of interest rate propagation.

Our work is also directly related to the literature investigating the heterogeneous effects of monetary policy conditional on financial and credit conditions, and debt or collateral values (see e.g., Hubrich and Tetlow (2015), Harding and Klein (2019), Aikman et al. (2019), Alpanda and Zubairy (2019), Ottonello and Winberry (2018), Beraja et al. (2019) and Cloyne et al. (2019)) and to the transmission of monetary policy through house prices (see e.g., Attanasio et al. (2011), Aoki et al. (2004), Campbell et al. (2012)).

Our empirical framework is motivated by insights provided by a large literature on the non-linear effects of monetary policy. The contributions of Weise (1999), Garcia and Schaller (2002), Lo and Piger (2005), Santoro et al. (2014), Caggiano et al. (2017) focus on the dependency of the effect of monetary policy to the state of the economy. Angrist et al. (2018) compare the real effects of tightening versus expansionary policies while Barnichon and Matthes (2014) and Tenreyro and Thwaites (2016) combine both types of non-linearities to analyse the transmission of monetary policy. Tillmann (2017) and Hubert (2018) also examine the state-dependence of monetary policy to uncertainty and central bank information shocks.

The rest of this paper is organised as follows. Section 2 sets out how we use our micro-data to measure borrowing constraints at the macro-level and construct our state-variable. Section 3 describes the identification of monetary shocks. Section 4 details the empirical strategy and the main result. Section 5 further explores what we can learn from our state-dependence results and Section 6 provides some robustness tests. Section 7 concludes.

## 2 Measuring Borrowing Constraints

### 2.1 Debt in the transmission mechanism

We exploit proprietary loan-level data to construct a clean measure of the economy’s exposure to highly indebted households, that we can use to investigate the transmission of monetary policy. Although monetary policy affects households along various dimensions, we focus on its transmission through the distribution of household mortgage debt. This enables us to capture the set of households with sizable amounts of illiquid wealth, which leaves them close to their budget constraint.

The income channel of monetary policy captures the reallocation and general-equilibrium effects of interest rate changes on household resources and behaviour. Lower interest rates lead to a redistribution of net income from savers to borrowers. The latter are more likely to

have high marginal propensities to consume so this redistribution can have aggregate effects (e.g., see Auclert (2019)). Beyond direct cash-flow effects, the aggregate impact of interest rates is likely to affect borrowers and savers via support to other real income flows associated with higher economic activity and a boost to labour income (e.g., see Kaplan et al. (2018)).

Variation in household income gearing refracts the impact of changes to interest rates because it alters the proportionate response of household disposable income. Those spending a large fraction of their post-tax income on mortgage repayments (i.e. those closer to their budget constraints) will therefore be more sensitive to the direct and indirect effects of monetary policy. Since behavioural responses are likely to be convex as households approach their budget constraints, a high fraction of mortgagors on high-LTI mortgages is likely to reflect a household sector that is significantly more responsive to monetary policy shocks.

The challenge for researchers - and policymakers - is to define what is meant by a *high*-LTI mortgage. By some measures, households in the UK are as indebted today as they ever have been because of the steep rise in borrowing relative to income over the last two decades. The mean LTI at origination has increased from around 2.5 in 2005 to just over 3 in 2018. During the same period, the 95<sup>th</sup> percentile LTI-at-origination increased from around 4 to around 4.7. To the extent debt magnifies the effect of monetary policy, that might imply monetary policy is as effective now as it ever has been - at least for those that have a mortgage.

We argue these simple indebtedness metrics miss some of the crucial elements of household debt decisions. Structural declines in long-term interest rates, regulatory changes to the banking system and subtle preference shifts are all likely to perturb the underlying level of debt that households are willing and able to take on. Historically, it has been difficult to construct household debt measures that properly account for the slower-moving developments and various structural breaks in the first two decades of the century. In order to define high-LTI mortgages we therefore appeal to micro data to control for household, banking and other structural characteristics across time and space.

## 2.2 Using PSD data to measure constraints

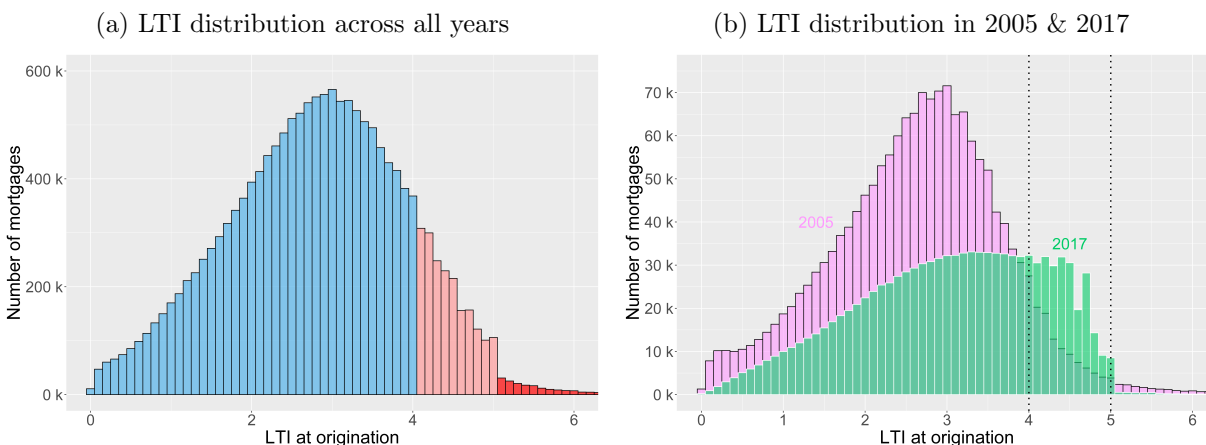
The mortgage data used in this study are derived from the universe of mortgage originations collected by the Financial Conduct Authority (FCA). This regulatory loan-level data set, known as the Product Sales Database (PSD), contains a wealth of information on the characteristics of all residential mortgages issued by lenders since April 2005. The PSD contains accurate information on borrower, loan and property characteristics associated with over 10 million residential mortgages issued between 2005 and 2017. Of those, around half are mortgages used to purchase a new house and the other half are refinancing transactions.

We use this detailed micro data to construct accurate measures of the distribution of mortgages during the period we study. In the first instance we track the proportion of mortgages issued above two threshold-LTI values every month between April 2005 and December 2017. The cut-offs of 4 and 5 constitute our initial definition of a *high*-LTI mortgage. The left-hand panel of Figure 1 shows that, across the whole sample, around 15% of mortgages

were issued with a LTI above 4 and less than 2% with an LTI above 5.

These LTI distributions have evolved over time and the full-sample panel masks some of the striking developments in UK mortgage issuance over the last two decades. The right-hand panel of Figure 1 shows the LTI distributions for 2005 and 2017 on the same axis in pink and green, respectively. While the mean has clearly shifted to the right, there has also been a noticeable shift in the shape of the distribution as various regulatory and prudential measures have limited the number of mortgages issued at very high LTIs. The contrast between these two distributions drives our desire to construct a metric that captures the proximity of the economy to its borrowing constraint using as much information as possible, while allowing for the various changes that have happened over the sample.

Figure 1: LTI distribution 2005-2017



Note: The left-hand panel plots the overall distribution of Loan-to-Income (LTI) ratios over our sample of 11,288,910 observations from 2005 to 2017. The dark red bars correspond to the LTI cut-off of 5 and the lighter red bars correspond to the cut-off of 4. The right-hand panel plots the distribution of LTI for the year 2005 (pink) and 2017 (green).


The cut-offs we use for the proportions of high-LTI mortgages are driven by an effort to construct a measure that best captures the non-linear change in behaviour associated with being close to the household budget constraint, and also exhibits sufficient variation over time. A typical mortgage with an LTI of 4 over our sample required households to spend a little over 35% of their disposable income on mortgage repayments. A one percent increase in interest rates would have pushed that proportion to around 40% for the average household, which is often associated with an increased risk of payment difficulties (e.g., see [Bank of England \(2017\)](#)). We judge that the proportion of mortgages above an LTI of 4 does a good job at capturing the non-linear sensitivity of households to changes in interest rates. The higher threshold of 5 is well into the tails of mortgage issuance and few lenders in recent times have originated many mortgages at these levels, in part due to regulatory changes. In addition, the overall distribution of LTI in Figure 1 shows a clear discontinuity at this level.

One might think a more natural metric is to look directly at the payment-to-income ratio (PTI).<sup>2</sup> In this paper we focus on the LTIs for three reasons. First, PTI measures a relatively

<sup>2</sup>PTI is sometimes referred to as the debt-service ratio (DSR).

transitory form of constraint and is influenced by interest rates and mortgage structuring decisions. Importantly, the average mortgage term has increased from 20 years before interest rates fell in 2008 to over 24 years today. In addition, PTIs can mask heterogeneity in the time-profile of mortgage payments: lower interest rates early on in the contract make large debt burdens more manageable, but repayments can vary over time for multiple reasons. In contrast, the LTI does a much better job at reflecting how close the borrower is to a longer-term measure of their lifetime budget constraint. Second, LTI corresponds to the metric macroprudential policymakers have used to design policy. In June 2014, the Bank of England’s Financial Policy Committee (FPC) limited the proportion of mortgages issued above an LTI of 4.5 to 15% of any given bank’s mortgage issuance. Since this policy focused attention on LTIs and was a significant driver of the change in the debt profiles following implementation, it makes sense to use this as our central metric.<sup>3</sup> Finally, and related to policy design, the LTI metric is attractive as it suffers very little measurement error. There is often not enough information available to impute the exact PTI of all mortgages, especially at the beginning of our sample.

Figure 2: LTI time-series 2005-2017



Note: This figure shows the monthly time-series of the share of LTI above 4 (left panel) and above 5 (right panel) relative to all mortgage originations in each month in the Product Sales Database.

Based on the threshold values of 4 and 5, the evolution of the monthly proportion of new high-LTI mortgages should reflect the degree to which borrowing constraints bind in the economy. Figure 2 shows how the proportion of mortgages above these thresholds has evolved over time. Both panels show that fluctuations in credit conditions, the housing market and bank and borrower leverage preferences have had a significant impact on these metrics over our sample. In the left-hand panel, there is an upward trend that reflects lower interest rates expanding the envelope of affordability and higher house prices requiring larger debt burdens. The right-hand panel shows a structural break in 2015 where regulation started to make very-high-LTI mortgages less attractive to issue.

<sup>3</sup>This macroprudential intervention was framed as an insurance mechanism to prevent aggregate over-indebtedness (Bank of England (2014)).



One way to deal with these underlying trends and structural breaks is to construct a state variable that uses de-trended values of these proportions. While attractive in its simplicity, we argue that this might still miss the latent state of the economy. For example, longer term interest rates have fallen since 2005, which has led to a right-ward shift in the LTI distribution and an increase in the number of households with high LTI mortgages. It is nevertheless challenging to claim that deviations from the reduced-form upward trend in the left-hand panel of Figure 2 captures the relative proportion of households close to their borrowing constraint. In order to better disentangle these factors, we take advantage of our micro-data to produce a more precise measure of how the share of constrained households has varied across time.<sup>4</sup>

## 2.3 Exploiting micro data to measure constraints

We exploit the detailed information we have on borrowers, properties and contracts to better identify fluctuations in the underlying indebtedness of households. To the best of our knowledge, this is the first study to construct an aggregate measure of high-indebtedness that controls for such a rich array of information. There are many factors that influence which houses people choose to live in and the types of mortgages people use to fund those choices. In order to understand the distribution of LTIs, we first need to take account of the contextual information that lies behind mortgage originations. At the heart of our micro-data approach is a loan-level regression of LTI (at origination) on a variety of pieces of information, shown in Equation 1.

$$LTI_i = \alpha + \alpha_{tm} + \alpha_l + \alpha_b + \beta X_i + \gamma MedLTI_{tq} + \epsilon_i \quad (1)$$

We employ time ( $\alpha_{tm}$ ), location ( $\alpha_l$ ) and bank ( $\alpha_b$ ) fixed effects to control for macro shocks, regional disparities and credit-supply drivers of LTI decisions. We also control for a vector of loan, borrower and property characteristics ( $X_i$ ) and we control for the median nationwide LTI ( $MedLTI_t$ ) for each quarter to capture developments in aggregate household credit markets, shown in Table 1. The main coefficients are followed by the categorical controls in the panel below. We control for mortgage-product choices, including how much to borrow, what kind of interest rate (essentially, fixed or adjustable-rate) and what proportion of their income to spend on mortgage servicing (the PTI ratio), imputed from other raw variables. Three regression specifications are shown in Table 1.

There are two general points worth making. First, the coefficients are estimated to a very high significance level, which is due in large part to the 11m observations that give a very clear read on what drives LTI decisions.<sup>5</sup> Second, the three specifications are very similar in the results they produce and the goodness of fit, they all have an r-squared of around 0.4. The specification we go on to employ for the analysis is in the third column.

In order to correct for when and where the mortgage was issued, we proceed in two steps. First, we compute the fitted values of this regression as a measure of the expected LTI of

---

<sup>4</sup>For robustness purposes, we explore the detrending strategy in the Appendix.

<sup>5</sup>We dropped observations with missing entries or data-entry errors.

Table 1: Loan-level regressions

	Loan-to-income ratio		
	(1)	(2)	(3)
Loan amount, £000,000s	2.1389*** (0.0033)	2.4434*** (0.0033)	1.5582*** (0.0033)
House price, £000,000s	-0.1753*** (0.0013)	-0.1516*** (0.0013)	0.1121*** (0.0013)
Interest rate, percent	-0.1210*** (0.0004)	-0.1195*** (0.0004)	-0.1170*** (0.0004)
PTI	2.2778*** (0.0017)	2.2219*** (0.0016)	2.0731*** (0.0016)
Age	-0.0157*** (0.0000)	-0.0154*** (0.0000)	-0.0091*** (0.0000)
Income, £0,000s	-0.0118*** (0.0000)	-0.0269*** (0.0000)	-0.0261*** (0.0000)
Income squared, £0,000s		0.0000*** (0.0000)	0.0000*** (0.0000)
LTV			0.0279*** (0.0000)
Median LTI	0.1897*** (0.0138)	0.1901*** (0.0132)	0.1862*** (0.0132)
Constant	2.3126*** (0.0395)	2.2883*** (0.0390)	0.9116*** (0.0377)
<b>Controls</b>			
Quarter	✓	✓	✓
Mortgage type	✓	✓	✓
Rate type	✓	✓	✓
Region	✓	✓	✓
Sole/joint income	✓	✓	✓
Originating bank	✓	✓	✓
LTV bucket	✓	✓	✓
Observations	11,288,910	11,288,910	11,288,910
Adjusted R2	0.39	0.40	0.44


Note: OLS estimation of Equation 1. \*\*p<0.05; \*\*\*p<0.01.  
Standard errors in parentheses.

each household at the point they got the mortgage. This steps enables us to control for the set of characteristics mentioned above. Since we want to compare the distribution of LTIs over time, we make one further adjustment that takes account of the trend of house prices and overall mortgage burdens. In the second step, we take the difference between the estimated partial effect of the median LTI at the end of 2017 and the estimated partial effect of the median LTI at the time the individual mortgage was issued. Since LTIs increased between 2005 and 2017 this essentially boosts the predicted LTI of early mortgages. We do the same exercise for the predicted values based on the coefficient on the loan amount just in case the numerator (loan value) was governed by a slightly different trend. The result is a set of conditional LTI values that counter-factually assumes each mortgage had been issued at the end of our sample. It gets us some way towards allowing us to compare mortgages on a like-for-like basis. Overall, the regression and prediction adjustments are a more pre-

cise way of understanding credit market outcomes relative to their underlying natural state.<sup>6</sup>

We calculate the monthly proportion of conditional LTIs above our main thresholds of 4 and smooth it by taking three-month rolling averages. The demeaned version of this series constitutes our state variable and is shown in Figure 3. The grey shading shows periods when the conditional share of high-LTI mortgages was above its sample average.

Figure 3: Share of conditional-LTI>4 and LTI>5 mortgages



Note: This figure shows the share of conditional LTI above 4 after the regression and prediction adjustments based on Equation 1 described above. For indicative purposes, the grey shaded areas correspond to periods when the LTI>4 state-variable is above its sample average. The LTI>5 state variable is also plotted although on a different scale.

Intuitively, our macro borrowing constraint measure take the information of a mortgage issued in Scotland in 2005 and allows us to compare what the LTI likely would have been for the same person in an average part of the UK in 2017. We are then able to calculate a monthly time-series for the proportion of like-for-like mortgages with an LTI above 4, all relative to the distribution of LTIs seen at the end of our sample. In months where more people than average took out high LTI mortgages (after all our adjustments) we believe more households were close to their borrowing constraints, and the economy as a whole was therefore more sensitive to shocks.

Our choice of using thresholds of conditional LTIs is our best attempt at identifying the underlying balance-sheet positions of households such that these constrained households may have a high marginal propensity to consume. As discussed above, our LTI metric captures the income channel, which can loosely be thought of a mixture of cash-flow and general equilibrium effects. Importantly, we care about these effects relative to the prevailing conditions in the economy. The interaction of this state variable and our monetary shock series will therefore shed some light on the state contingency of monetary policy.

<sup>6</sup>In section 6.2, we provide robustness tests using a simple de-trending of the raw monthly share of LTI above 4, or using the fitted values without the correction for the trend in median LTI and house prices.

### 3 The identification of monetary shocks

When investigating the causal impact of monetary policy on consumption, we face the usual endogeneity issue of the policy instrument. To identify exogenous monetary policy innovations, we use the [Romer and Romer \(2004\)](#) approach that consists of singling out the systematic endogenous part of monetary policy. This approach capitalises on three advantages: it is appropriate for low-frequency data and relatively short samples, it is parsimonious in the number of parameters to estimate, and it enables us to control specifically for the information set available to policymakers at the time of policy decisions. [Coibion \(2012\)](#) have shown that this identification strategy works well to investigate the transmission of monetary policy to the real economy and the [Romer and Romer \(2004\)](#) approach was first applied to UK data by [Cloyne and Huertgen \(2016\)](#).<sup>7</sup> Equation 2 represents a simplified version of the central bank reaction function, such that we can decompose the systematic endogenous response of policy to the economy and an exogenous component:

$$\Delta i_t = f(\Omega_t) + \epsilon_t^i \quad (2)$$

Changes in the policy instrument,  $\Delta i_t$  are regressed on the central bank information set,  $\Omega_t$ , at the time of the decision. The residual,  $\epsilon_t$ , represents the monetary shock series. [Blanchard et al. \(2013\)](#) and [Miranda-Agrippino and Ricco \(2017\)](#) have demonstrated how information frictions modify the econometric identification problem: in the presence of non-nested information sets, exogenous monetary innovations should also be made orthogonal to private agents' information set. [Ramey \(2016\)](#) and [Miranda-Agrippino and Ricco \(2017\)](#) have shown that standard monetary shocks like the one estimated in Equation 2 can be predicted with private information sets. We therefore augment this empirical model to include a proxy of private agents' information set.

There are at least two reasons why the conventional policy rate, Bank rate in the UK, may not be relevant to measure the stance of monetary policy over the sample period considered. First, partly because of the Zero Lower Bound (ZLB), monetary policy has become multi-dimensional since the Great Recession. The implementation of asset purchases, forward guidance and targeted liquidity provisions has rendered the policy rate insufficient to capture the overall stance of monetary policy. Second, a large consensus has formed about the content of monetary policy news: [Gürkaynak et al. \(2005\)](#), [Campbell et al. \(2012\)](#) and [Hanson and Stein \(2015\)](#) suggest that the main piece of information on central bank announcement days relates to changes in the future likely policy path (whether it is the policy rate during a period of conventional monetary policy, or asset purchases in the most recent period) over the next several quarters, as opposed to changes in the current policy stance. A simple and transparent way to encompass in one single variable the multi-dimensionality of monetary policy, and to capture the information about the expected path of policy over a given horizon, is to use nominal sovereign yield at this horizon as a proxy for the overall monetary policy stance. Following [Hanson and Stein \(2015\)](#), we use 2-year interest rates.<sup>8</sup>

---

<sup>7</sup>We assess the impact of this choice by also using monetary surprises estimated using intraday data.


<sup>8</sup>We also use [Krippner \(2013\)](#)'s shadow rate to assess the robustness of our main result to this choice.

Equation 2 can thus be rewritten as follows:

$$\begin{aligned}
\Delta i_t = & \alpha + \rho i_{t-1} + \sum_{j=1}^3 \lambda_j \mathbb{E}_t^{CB} \pi_{t+j} + \sum_{j=1}^3 \gamma_j \mathbb{E}_t^{CB} x_{t+j} \\
& + \sum_{j=1}^3 \phi_j \Delta \mathbb{E}_t^{CB} \pi_{t+j} + \sum_{j=1}^3 \varphi_j \Delta \mathbb{E}_t^{CB} x_{t+j} \\
& + \sum_{j=1}^3 \psi_k U_{t-k} + \kappa \mathbb{I}_t^{IR} + \eta_{sw} \mathbb{E}_t^{sw} \pi_{t+2} + \eta_{sef} \mathbb{E}_t^{sef} \pi_{t+2} + \epsilon_t^i
\end{aligned} \tag{3}$$

where  $i_t$  is the policy instrument, proxied by nominal 2-year interest rates,  $\mathbb{E}_t^{CB} \pi_{t+j}$  and  $\mathbb{E}_t^{CB} x_{t+j}$  are the Bank of England’s inflation and output forecasts at horizons 1, 2 and 3 years ahead.<sup>9</sup> Equation 3 also includes  $U_{t-k}$  the unemployment rate at different lags  $k$ .  $\mathbb{I}_t^{IR}$  is a dummy that takes the value 1 in months when the Inflation Report is published.<sup>10</sup>  $\mathbb{E}_t^{sw} \pi_{t+2}$  is a market measure of inflation expectations at the 2-year horizon based on inflation swaps corrected for liquidity premia, and  $\mathbb{E}_t^{sef} \pi_{t+2}$  is a measure of inflation expectations at the 2-year horizon from professional forecasters based on the *Survey of External Forecasters*.<sup>11</sup> The residuals  $\epsilon_t^i$  are the proxy for the monetary shocks.

Figure 4: Monetary policy shocks



Note: This figure shows our instrument for monetary shocks. The time series  $\epsilon_t^i$ , estimated with equation 3, refers to unpredictable changes in 2-year nominal rates.

Figure 4 plots this estimated monetary shock series. As expected, the largest innovations occur around 2008 and 2009 when the Global Financial Crisis was hitting the UK economy. We have also tested that these monetary shocks are unpredictable from movements in macroeconomic data.<sup>12</sup>

<sup>9</sup>The use of forecasts to measure policymakers and private agents’ information sets relates to three points. First, forecasts are real-time data, such that they enable to make sure that the information set available at the time of the decision. Second, forecasts capture the forward-looking characteristics of information sets. Third, forecasts encompass rich information sets and work as a FAVAR model (see Bernanke et al. (2005)) as they summarize a large variety of variables.

<sup>10</sup>Policymakers may be more likely to update policy when they update their macroeconomic projections.

<sup>11</sup>We consider two different types of inflation expectations, from market participants and from professional forecasters since there may be non-nested information sets among private agents as well. We focus on inflation expectations, consistent with the remit of the Bank of England to target a 2% inflation rate.

<sup>12</sup>These tests are available from the authors upon request.

## 4 The role of constrained households in the transmission of monetary policy


We aim to measure the dynamic effect of monetary shocks conditional on the share of households close to their borrowing constraint. To do so, we use local projections as proposed by Jordà (2005) with our externally identified instruments for monetary policy. We interact the monetary shock series with our state-variable  $S_t$  described in Section 2. Local projections have become a popular tool to compute impulse responses since they are robust to model misspecification. Impulse response functions obtained from VARs may inadvertently impose excessive restrictions on the endogenous dynamics, while the local projection method is more flexible and allows for non-linearities. The Jordà (2005) method requires estimating a series of  $k$  regressions for each horizon, with the estimated coefficient representing the response of the dependent variable at the horizon  $k$  to a given exogenous shock at time  $t$ . It is close in spirit to an autoregressive distributed lag model or to a smooth transition-local projection model. Equation 4 is estimated  $k$  times as follows:

$$\begin{aligned}
 y_{t+k} = & \alpha_k + \beta_k \epsilon_t^i + \gamma_k S_t + \delta_k S_t \cdot \epsilon_t^i \\
 & + \Theta_{k,j} \sum_{j=1}^3 (\epsilon_{t-j}^i + S_{t-j} + S_{t-j} \cdot \epsilon_{t-j}^i) \\
 & + \lambda_k X_t + \phi_k S_t \cdot X_t + \epsilon_t
 \end{aligned} \tag{4}$$

$y_{t+k}$  is a measure of consumption at different horizons  $t+k$ ,  $\epsilon_t^i$  the monetary shock,  $S_t$  the state-variable capturing the share of highly indebted households, and  $S_t \cdot \epsilon_t^i$  is their interaction.  $S_t$  is a continuous variable, and we show the effect of monetary shocks for two different values of  $S_t$  (mean  $\pm$  1 SD). We control for the effects of three lags of the monetary shock, the state-variable and their interaction.  $X_t$  is a vector of controls related to macroeconomic dynamics, the structure of the mortgage and housing markets, and LTI dynamics over our sample. This vector is also interacted with the state variable. Equation 4 is estimated using OLS from April 2005 to December 2017. We compute heteroskedasticity robust standard errors. Testing the hypothesis that borrowing constraints matter for the transmission of monetary policy is tantamount to testing that  $\delta_k$  is significantly different from zero.

$X_t$  includes one lag of the dependent variable as well as the contemporaneous value and one lag of the following variables. First, we include the annual change in employment, which captures the state of the economy, and in wages to control for changes in household income that could explain changes in consumption. In addition,  $X_t$  includes the average mortgage duration, the share of newly-issued floating-rate mortgages and the annual change in the share of floating mortgages in the overall stock so as to control for changes in the structure of the mortgage market over our sample. Because credit conditions are endogenous to the monetary transmission mechanism, we include the average interest rate of new mortgages and the flow of (secured and unsecured) lending in order to control for price and quantity dynamics of the overall credit market. Finally, we also control for the share of new mortgages that come from remortgages, median LTI and house prices to make sure we are not capturing changes in household preferences and the housing market over time.

Figure 5: Monetary Transmission in the Presence of High-LTI Mortgages



Note: This figure shows the estimates of the effect of  $\epsilon_t^i$  over 12 months for consumption, based on the OLS estimation of Equation 4 over the sample April 2005 - December 2017. The black empty circles correspond to the effect of a contractionary monetary policy with a high share (1 SD above the sample mean) of households above a conditional LTI of 4. The blue full circles correspond to the effect of a contractionary monetary policy when there is a low share (1 SD below the sample mean). The shaded areas represent 68 and 90% confidence intervals.


In a standard treatment of policy transmission, aggregate credit conditions have been shown to have some effect on the behaviour of the economy. In this paper, we are particularly interested in the distribution of indebtedness precisely because it might tell us about how monetary policy can act in non-linear ways. Highly indebted households close to their borrowing constraints are likely to have higher marginal propensities to consume, but are also mechanically more exposed to changes in debt servicing costs. In the following analysis, we focus on the response of non-durable, durable and total consumption in response to a monetary contraction to test our central hypothesis. If our state variable does not matter, we would expect overlapping trajectories for consumption responses when the share of constrained households is high and low.

We find that the transmission of monetary policy on all forms of consumption does depend on the share of constrained households. This is consistent with the literature demonstrating the key role of households with high levels of illiquid wealth. As shown in the first panel of Figure 5, a one percentage point monetary shock policy shock, conditional on a high share of constrained households, leads to a peak impact on dampened non-durable consumption of around 4pp after eight quarters in the grey swathe. The effect on durable consumption in the middle panel is more modest and the peak impact occurs after six months. The impact on total consumption is therefore close to 5pp just before six months. In contrast, the effect of monetary policy on all forms of consumption when the share of highly-indebted households is low, shown in the blue swathe, is insignificantly different from zero across all horizons and statistically different from the grey swathe for much of the projections.

The gaps between the grey and blue swathes provide evidence that the effectiveness of monetary policy might vary according to the share of borrowers close to their financial constraint. In order to explore this further, we modify our state variable in different dimensions to make sure we are doing a good job at capturing states of the economy when households are more or less constrained. Figure 6 shows very similar consumption responses when we

look at a higher conditional LTI threshold of 5.<sup>13</sup> In the Appendix, we show that we find similar consumption responses when you restrict our main state variable to subsets of the population that might intuitively be thought of as being constrained.

Figure 6: Monetary Transmission conditional on LTI>5 mortgages



Note: This figure shows the estimates of the effect of  $\epsilon_t^c$  over 12 months for consumption, based on the OLS estimation of Equation 4 over the sample April 2005 - December 2017. The black empty circles correspond to the effect of a contractionary monetary policy with a high share (1 SD above the sample mean) of households above a conditional LTI of 5. The blue full circles correspond to the effect of a contractionary monetary policy when this share is low (1 SD below the sample mean). The shaded areas represent 68 and 90% confidence intervals.

The blue swathes in Figure 6 show counter-intuitive responses of a restrictive monetary shock when the share of high-LTI mortgages is low. We see at least three plausible explanations for these counter-intuitive responses. First, if monetary policy is truly state contingent and systematic policy responses are calibrated to a period of average credit conditions, then monetary shocks might not be enough to prevent the economy from over shooting when the economy is less responsive to changes in interest rate. Second, when monetary policy has less direct traction on agents' behaviour, the signalling channel of monetary policy might become more important. Policy decisions revealing policymakers' views on the macroeconomic outlook would be consistent with the blue swathes. Third, a neo-fisherian mechanism might be at work such that inflation, investment and consumption increase after a tightening.

Overall, our results suggest monetary policy might be more effective in the presence of a large number of households close to their borrowing constraints. In the next section, we explore the possible channels through which this effect might be working.

## 5 Exploring further the state-dependent effect

### 5.1 Uncovering the transmission channels


In the previous section, we provide evidence that consumption is more responsive to monetary policy when there is a large share of households close to their borrowing constraints. This section now investigates which mechanisms might be responsible for this state-dependent transmission of monetary policy. Before turning to that, Figure 7 presents some evidence

<sup>13</sup>We find a similar result with a threshold at 4.5 consistent with the FPC limit.



for the asymmetric effects of monetary policy. We simply augment Equation 4 with a triple interaction term for when policy is expansionary or restrictive. The left-hand panel shows that restrictive monetary policy has a significantly larger impact on consumption than expansionary policy, in line with Tenreyro and Thwaites (2016)). This result is consistent with restrictive monetary policy moving households closer to their budget constraint, while expansionary monetary policy leaves households on their Euler equation. Since the effect of monetary policy is found to be asymmetric, it is important to examine whether the main result shown in Figure 5 is not driven by a confluence of factors and a composition effect: there might be more restrictive monetary shocks during periods when borrowing constraints are tighter. The right-hand panel of Figure 7 plots consumption responses to restrictive shocks (identified specifically in a non-linear setting) when borrowing constraints are high or not. We find that the main result that monetary policy has more traction still holds.

Figure 7: Disentangling restrictive and expansionary policy decisions



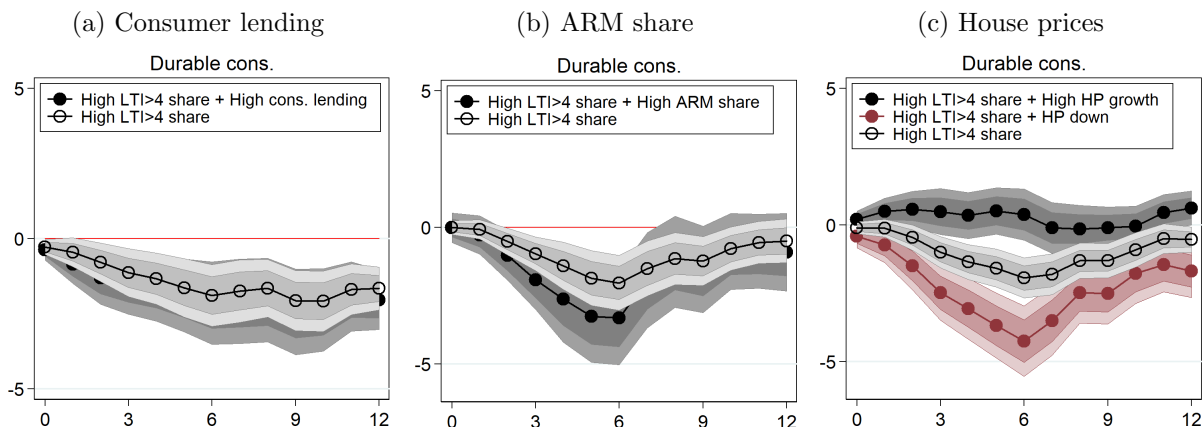
Note: This figure shows the estimates of the effect of  $\epsilon_t^i$  over 12 months for consumption, based on the OLS estimation of an augmented Equation 4 over the sample April 2005 - December 2017, conditional on a the share of households above a conditional LTI of 4 and on whether monetary policy is restrictive or expansionary. The right panel shows the effect of expansionary (filled circles) and restrictive (empty circles) monetary shocks when the share of LTI above 4 is high. The left panel shows the effect of restrictive monetary shocks when the share of LTI above 4 is high (black empty circles) and low (blue filled circles). The shaded areas represent 68 and 90% confidence intervals.

There are several channels that might explain our findings that monetary policy is more effective when a large fraction of people are highly indebted. Inter-temporal substitution, where higher interest rates make it more expensive to reallocate resources across time, is one natural candidate. But when the share of households close to their budget constraint (so off their Euler equation) is high, the effect of monetary policy through inter-temporal substitution should actually be dampened. To test whether the drop in consumption is driven by people bringing forward less spending than they otherwise would have, we look at how the interplay of consumer lending and monetary policy in driving consumption. The first panel of Figure 8 shows the result of a triple interaction of the monetary policy shock, the high state variable and a secondary state variable that represents periods with an above-average flow of consumer lending. If inter-temporal substitution was playing a significant role, we would expect to see durable consumption falling more after a monetary contraction in the

presence of high volumes of consumer lending. In fact, both swathes are somewhat overlapping, which suggests this channel might not be particularly important.

Another possibility is that highly indebted households react to cash-flow shocks that stem from changes in interest rates. For this mechanism to play a role, the increase in mortgage servicing costs leads to a fall in disposable income and spending (perhaps because households do little to smooth consumption successfully over time). To test this channel we perform a triple interaction of the monetary shock, our high state variable and a secondary state variable that represents a high share of adjustable-rate mortgages (ARMs) in originations. The idea behind this experiment is that we would expect the drop in consumption to be more dramatic when the pass-through of interest rates is quicker to mortgage payments. We find weak evidence of the cash-flow channel in the middle panel of Figure 8, since the point estimates are only just significantly different at the 90% level.

Figure 8: Transmission channels - Hypothesis testing



Note: This figure shows the estimates of the effect of  $\epsilon_t^i$  over 12 months for consumption, based on the OLS estimation of augmented Equation 4 over the sample April 2005 - December 2017, conditional on a high share of households above a conditional LTI of 4 and a secondary state-variable. The black empty circles show the effect of a restrictive monetary policy when the secondary state-variable is at its sample mean. The black filled circles corresponds to the effect of a restrictive monetary policy when the secondary state-variable is 1 S.D. above its sample average. The red circles show the effect of a restrictive monetary policy when house prices are decreasing. The shaded areas represent 68 and 90% confidence intervals.


Household consumption choices are affected by cash-flows but also their balance sheet positions. One channel that has attracted a lot of recent attention is the collateral channel of monetary policy, where interest rate changes amplify behaviour via their effect on asset prices. For households this comes in two main flavours. First, an increase in asset prices makes it easier for households to refinance their mortgages and in so doing reduce their overall servicing costs.<sup>14</sup> Second, higher house prices make households wealthier by increasing their net asset position and reduce the need for precautionary or retirement saving. With both of these mechanisms in mind, we perform a triple interaction of our monetary shock, our high state variable and a secondary state variable reflecting above average house price growth.

<sup>14</sup>In general, UK mortgage interest rates are an increasing function of the LTV, so changes in asset prices can reduce the interest rate on the outstanding balance. Households might also take out equity but we control for this behaviour in the local projection.

The right-hand panel of Figure 8 suggests that the refinancing or wealth effect of house prices might be significant and implies that the restrictive effect of monetary policy might be at least partially offset in an environment of recent high house price growth (filled black dots). The red swathe suggests restrictive monetary policy is especially potent in an environment of falling house prices. This result is consistent with some recent evidence on the refinancing channel of monetary policy (see Eichenbaum et al. (2018) and Beraja et al. (2019)) and that consumption responses depend on debt and house prices (Berger et al. (2017)).

## 5.2 Assessing the propagation to the economy

The less direct effects of monetary policy work through its influence on overall aggregate income and activity. Although somewhat less studied, there are good reasons to think that these general-equilibrium effects could play a substantial role in the overall traction of monetary policy (e.g., see Kaplan et al. (2018)). To explore this mechanism we take our state contingent analysis and move beyond an examination of consumption. Figure 9 shows four panels that give us some clues about how monetary policy differentially affects the economy when more households are close to their budget constraint.




Note: This figure shows the estimates of the effect of  $\epsilon_t^i$  over 12 months for different macro variables, based on the OLS estimation of equation 4 over the sample April 2005 - December 2017. The black empty circle corresponds to the effect of a restrictive monetary policy with a high share of LTI above 4 whereas the blue full circle corresponds to the effect of a restrictive monetary policy with a low share of LTIs above 4. The shaded areas represent 68 and 90% confidence intervals.

The first row suggests that a restrictive monetary policy shock has a relatively limited impact on the labour market. There is some tentative evidence that wages react slightly more strongly in the high state but the employment swathes are very similar. The bottom row does, however, suggest that overall economic activity (either measured by GDP growth or the more frequent industrial production series) is significantly more responsive to monetary policy when the economy is closer to its borrowing constraint. One potential reason for this would be a knock-on effect of lower spending on firm behaviour.

### 5.3 The role of heterogeneity

One important aspect of this paper is to focus on the distribution issue of households' borrowing constraints. Some recent works (Cloyne et al. (2019) or Gelos et al. (2019)) have shown the importance of household debt, overall. In order to make sure that the differentiated effects that we evidence are not simply linked to more indebtedness households, we have included the median LTI of each month in equation 4. To further ensure that our borrowing constraint state-variable is not capturing the variation in average income gearing over our sample, we re-estimate equation 4 with the monthly median LTI or mean LTI as the state-variable. As a complementary test that it is the distribution of borrowing constraints more than the average that matters, we assess the state-dependence effect of monetary policy conditional on the monthly mean LTV.

Figure 10: State-dependent effect of monetary policy conditional on average metrics



Note: This figure shows the estimates of the effect of  $\epsilon_t^i$  over 12 months for different macroeconomic variables, based on the OLS estimation of equation 4 over the sample April 2005 - December 2017. The red empty circle corresponds to the effect of a restrictive monetary policy with a high level of the median LTI (mean LTI or mean LTV) in each month whereas the yellow full circle corresponds to the effect of a restrictive monetary policy with a low level of the median LTI (mean LTI or mean LTV) in each month. The shaded areas represent 68 and 90% confidence intervals.

Figure 10 shows the state-dependent effect of monetary policy when interacted with such state-variables. In all three cases, the profiles of the impulse responses of consumption are different from the ones from our baseline specification. The main outcome is that they are not statistically different one from the other such that monetary policy has similar effects on consumption irrespective of the regime in which average borrowing constraints are.

## 6 Extensions

### 6.1 The impact of the Great Recession

In our central specification, we use the presence of high-LTI mortgages to proxy for borrowing constraints that might lead to changes in households' behaviour. But the fact that our data start in 2005 and therefore only span one credit cycle might be cause for concern, especially since monetary policy has been at the zero lower bound since 2009. One way to ensure our main state-variable is indeed capturing our intended channels is to introduce a placebo state-variable in the form of a dummy variable that takes the value of one from October 2008 and zero beforehand.

The left-panel of Figure 11 shows the impulse responses when we split the sample into these two periods. The red swathe this time captures the transmission of monetary policy after the Financial Crisis in October 2008. If our credit conditions state variables are in fact only capturing the structural break in the economy after the collapse of Lehman Brothers we would expect to see similar impulse responses in this placebo experiment. The main takeaway from Figure 11 is that monetary policy has had slightly more traction in the period before 2008, but this does not explain the profiles we observe in Figure 5. Indeed, the consumption responses are not statistically different one from the other. These profiles suggest that the effect we are capturing in the main specification is not related to the Great Recession and the subsequent low policy rate regime.

In addition, we test whether the main result stems from the Great Recession period specifically. During this period, the freezing of the interbank market and the credit crunch could have impacted households behaviour and how they respond to monetary policy. We therefore run our main specification when excluding the 12 months that followed the bankruptcy of Lehman Brothers, period during which the VIX was at exceptional high levels and GDP growth was negative. The right panel of Figure 11 shows the impulse responses when we exclude this particular sample. The profiles of the main result that we observe in Figure 5 hold: monetary policy has more effect when the share of households with strong borrowing constraints is high.

### 6.2 Sensitivity analysis about borrowing constraints

Another area of our results that is worth examining is the definition of the thresholds for the state-variables we use. The main assumption of this paper is that the chosen levels for the LTI threshold (4 and 5) represent an economy close to its borrowing constraints. We also provide impulse responses for a state variable computed with a threshold at 4.5 that makes sense in terms of the context of mortgage market practices and regulation.

One concern with the share of mortgages with a LTI above a certain threshold could be that it would capture households with strong preferences for debt more than households close to their income constraints. For instance, households with very high revenues and low marginal propensities to consume could have a very high LTI but still a comfortable

disposable income after mortgage payments. We therefore compute the share of mortgages when (i) LTI is above 4 and PTI is above 30%, (ii) LTI is above 4 and LTV is above 90%, and (iii) LTI is above 4 and LTV is above 75%. The idea is that these households close to these constraints are likely to be effectively close to their budget constraint. We look at two different levels (90 and 75%) of the LTV condition to make sure that we are not confounding factors and giving too much weight to the collateral channel (when using LTV above 90%) while we aim to capture households close to their budget constraints. Another way to alleviate this potential concern is to compute the share of LTI above 4 while excluding high-income households (the upper 25% of the distribution) or the share of LTI above 4 when focusing on older people (the upper 50% of the distribution) whose income are lower and the profile is likely not to increase. Figure 12 shows that we get a similar state-dependent transmission of monetary policy to consumption using these various variations on these thresholds.

In order to further assess the robustness of our results to the identification of borrowing constraints, we compute two alternative specifications of our state variable. First, we use the fitted value of Equation 1 directly without correcting for the positive trend in median LTI and house prices across our sample to compute the conditional share of LTI above 4. The identification of borrowing constraints therefore relies exclusively on the characteristics we control for in Equation 1. Second, we completely abstract from Equation 1 and compute the share of raw LTI above 4. We de-trend (using the one-sided Christiano-Fitzgerald filter) that share to control for changes in the mortgage and housing markets and in monetary and regulatory policies. The underlying (strong) assumption is that all these changes are captured by the trend. Figure 13 shows that we get a similar state-dependent transmission of monetary policy using these two alternative ways of representing borrowing constraints.

### 6.3 Alternative monetary policy identifications

Following an abundant literature, our central results rely on well-identified monetary shocks being used to estimate the causal effect of monetary policy on various economic variables. Although there is good justification to use the surprise element of the two-year risk-free nominal yield, it is important to demonstrate that we obtain a similar picture using other methodologies.

One reason to use the change in the two-year rate is that monetary policy was pushing up against the zero lower bound during much of our sample, so there was little change in the main policy rate. An alternative specification uses monetary shocks estimated off the change in Bank rate before 2009 and changes in the shadow rate thereafter, as calculated by Krippner (2013). This approach has the advantage of directly incorporating the multidimensional aspects of unconventional monetary policies, such as the forward guidance policy or developments in the Bank of England’s asset purchases during the Great Recession. Although conceptually this is a different measure of the stance of monetary policy, Figure 15 shows consistent results.

The same pattern of results also holds when we vary the conditioning assumptions used in the construction of the exogenous monetary innovation. Another abundant literature uses

high-frequency event-study methodology to measure monetary surprises. These are shocks to the information set of private agents. We follow [Hanson and Stein \(2015\)](#) to compute these monetary surprises. Although the confidence intervals at the 90% level overlap in the right panel of [Figure 15](#) - this might be due to the frequency mismatch between financial market responses and consumption responses -, the pattern of consumption responses to monetary policy when households are close to their borrowing constraint or not is consistent with our previous results.

## 7 Conclusion

Our findings provide empirical evidence showing how changes in monetary policy feed through to different parts of the economy via households' borrowing constraints. Using loan-level data enables us to focus on these constrained households holding illiquid wealth that have received much attention in the theoretical literature. Importantly, the efficacy of monetary policy appears to be much more affected by the distribution of debt rather than the overall indebtedness of the household sector. Although monetary policy has long and variable lags, we show that the distribution of household indebtedness might explain some of the variation in the potency of monetary policy over the business cycle.

We can use our state variable to explore some of the plausible channels through which monetary policy might operate. On the whole, our results are intuitive and shed some light on where further work might fruitfully make progress. In particular, the interaction of debt heterogeneity and movements in regional and national house prices seems like an obvious avenue to explore in future work. We also provide some tentative evidence that changes in monetary policy might lead to spillovers from consumption to firm behaviour, which is consistent with the relatively large role of *indirect effects* in many heterogeneous agent models.

Even if monetary policy makers can influence loan pricing, they have few tools to influence the distribution of debt in the economy. On the other hand, macroprudential policymakers can do exactly this with policies directed at households and lenders. Our results suggest that there is an important interaction between monetary and macroprudential policy, and therefore a potential role for co-ordination across policy committees.

## References

- Aikman, D., Lehnert, A., Liang, N., and Modugno, M. (2019). Credit, financial conditions, and monetary policy transmission. *International Journal of Central Banking*, *forthcoming*.
- Alpanda, S. and Zubairy, S. (2019). Household debt overhang and transmission of monetary policy. *Journal of Money, Credit and Banking*, 51(5):1265–1307.
- Angrist, J., Jordà, Ò., and Kuersteiner, G. (2018). Semiparametric estimates of monetary policy effects: string theory revisited. *Journal of Business and Economic Statistics*, 36(3):371–387.

- Aoki, K., Proudman, J., and Vlieghe, G. (2004). House prices, consumption, and monetary policy: a financial accelerator approach. *Journal of Financial Intermediation*, 13(4):414–435.
- Attanasio, O., Banks, J., and Tanner, S. (2002). Asset holding and consumption volatility. *Journal of Political Economy*, 110(4):771–792.
- Attanasio, O., Leicester, A., and Wakefield, M. (2011). Do house prices drive consumption growth? The coincident cycles of house prices and consumption in the UK. *Journal of the European Economic Association*, 9(3):399–435.
- Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–67.
- Bank of England (2014). Financial Stability Report. Issue No. 35, June 2014.
- Bank of England (2017). Financial Stability Report. Issue No. 41, June 2017.
- Barnichon, R. and Matthes, C. (2014). Measuring the non-linear effects of monetary policy. *Federal Reserve Bank of Richmond Manuscript*.
- Beraja, M., Fuster, A., Hurst, E., and Vavra, J. (2019). Regional heterogeneity and the refinancing channel of monetary policy. *Quarterly Journal of Economics*, 134(1):109–183.
- Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2017). House prices and consumer spending. *Review of Economic Studies*, 85(3):1502–1542.
- Bernanke, B., Boivin, J., and Elias, P. (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (favar) approach. *Quarterly Journal of Economics*, 120(1):387–422.
- Bilbiie, F. and Ragot, X. (2017). Optimal monetary policy and liquidity with heterogeneous households. CEPR Discussion Paper 11814.
- Blanchard, O., L’Huillier, J.-P., and Lorenzoni, G. (2013). News, noise, and fluctuations: An empirical exploration. *American Economic Review*, 103(7):3045–70.
- Bunn, P., Le Roux, J., Reinold, K., and Surico, P. (2018). The consumption response to positive and negative income shocks. *Journal of Monetary Economics*, 96:1–15.
- Caggiano, G., Castelnuovo, E., and Nodari, G. (2017). Uncertainty and monetary policy in good and bad times. RBA Working Paper 2017-06.
- Calza, A., Monacelli, T., and Stracca, L. (2013). Housing finance and monetary policy. *Journal of the European Economic Association*, 11(1):101–122.
- Campbell, J., Evans, C., Fisher, J., and Justiniano, A. (2012). Macroeconomic effects of federal reserve forward guidance. *Brookings Papers on Economic Activity*, 2012:1–80.



- Carroll, C., Slacalek, J., Tokuoka, K., and White, M. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020.
- Cloyne, J., Ferreira, C., and Surico, P. (2019). Monetary policy when households have debt: new evidence on the transmission mechanism. *Review of Economic Studies*, forthcoming.
- Cloyne, J. and Huertgen, P. (2016). The macroeconomic effects of monetary policy: A new measure for the United Kingdom. *American Economic Journal: Macroeconomics*, 8(4):75–102.
- Coibion, O. (2012). Are the Effects of Monetary Policy Shocks Big or Small? *American Economic Journal: Macroeconomics*, 4(2):1–32.
- Cumming, F. (2018). Mortgages, cash-flow shocks and local employment. Bank of England Working Paper 773.
- Debortoli, D. and Galí, J. (2017). Monetary policy with heterogeneous agents: Insights from tank models. *Manuscript*.
- Di Maggio, M., Kermani, A., Keys, B., Piskorski, T., Ramcharan, R., Seru, A., and Yao, V. (2017). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review*, 107(11):3550–88.
- Eggertsson, G. and Krugman, P. (2012). Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach. *Quarterly Journal of Economics*, 127(3):1469–1513.
- Eichenbaum, M., Rebelo, S., and Wong, A. (2018). State dependent effects of monetary policy: The refinancing channel. NBER Working Paper 25152.
- Finck, D., Schmidt, J., and Tillmann, P. (2018). Mortgage debt and time-varying monetary policy transmission. MAGKS Papers on Economics 2018-09.
- Flodén, M., Kilström, M., Sigurdsson, J., and Vestman, R. (2017). Household debt and monetary policy: Revealing the cash-flow channel. CEPR Discussion Paper 12270.
- Garcia, R. and Schaller, H. (2002). Are the effects of monetary policy asymmetric? *Economic Inquiry*, 40(1):102–119.
- Gelos, G., Grinberg, F., Mancini-Griffoli, T., Khan, S., Narita, M., and Rawat, U. (2019). Has higher household indebtedness weakened monetary policy transmission? IMF Working Paper 2019-11.
- Georgarakos, D. and Tatsiramos, K. (2019). Monetary policy transmission to consumer financial stress and durable consumption. *SSRN 3386611*.
- Greenwald, D. (2018). The mortgage credit channel of macroeconomic transmission. MIT Sloan Research Paper 5184-16.

- Gürkaynak, R., Sack, B., and Swanson, E. (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1):55–93.
- Hanson, S. and Stein, J. (2015). Monetary policy and long-term real rates. *Journal of Financial Economics*, 115(3):429–448.
- Harding, M. and Klein, M. (2019). Monetary policy and household deleveraging. Technical report, DIW Discussion Papers.
- Hubert, P. (2018). State-dependent effects of monetary policy: the central bank information channel. *Manuscript*.
- Hubrich, K. and Tetlow, R. (2015). Financial stress and economic dynamics: The transmission of crises. *Journal of Monetary Economics*, 70:100–115.
- Iacoviello, M. (2005). House prices, borrowing constraints, and monetary policy in the business cycle. *American Economic Review*, 95(3):739–764.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.
- Kaplan, G., Moll, B., and Violante, G. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743.
- Kaplan, G. and Violante, G. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.
- Kaplan, G., Violante, G., and Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, 2014(1):77–138.
- Krippner, L. (2013). Measuring the stance of monetary policy in zero lower bound environments. *Economics Letters*, 118(1):135–138.
- Lo, M. and Piger, J. (2005). Is the response of output to monetary policy asymmetric? evidence from a regime-switching coefficients model. *Journal of Money, Credit and Banking*, 37(5):865–886.
- Mian, A. and Sufi, A. (2011). House prices, home equity-based borrowing, and the US household leverage crisis. *American Economic Review*, 101(5):2132–56.
- Miranda-Agrippino, S. and Ricco, G. (2017). The transmission of monetary policy shocks. OFCE Working Paper 2017-15.
- Misra, K. and Surico, P. (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, 6(4):84–106.
- Ottoneo, P. and Winberry, T. (2018). Financial heterogeneity and the investment channel of monetary policy. NBER Working Paper 24221.

- Piskorski, T. and Seru, A. (2018). Mortgage market design: Lessons from the great recession. *Brookings Papers on Economic Activity*, 2018(1):429–513.
- Ragot, X. (2014). The case for a financial approach to money demand. *Journal of Monetary Economics*, 62:94–107.
- Ramey, V. (2016). Macroeconomic shocks and their propagation. In *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier.
- Romer, C. and Romer, D. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084.
- Santoro, E., Petrella, I., Pfajfar, D., and Gaffeo, E. (2014). Loss aversion and the asymmetric transmission of monetary policy. *Journal of Monetary Economics*, 68:19–36.
- Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: US monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74.
- Tillmann, P. (2017). Monetary policy uncertainty and the response of the yield curve to policy shocks. Joint discussion paper series in economics.
- Weise, C. (1999). The asymmetric effects of monetary policy: A nonlinear vector autoregression approach. *Journal of Money, Credit and Banking*, 31(1):85–108.