The Role of Borrowing Constraints
in the Transmission of Monetary Policy*

Fergus Cumming[†]   Paul Hubert[‡]

February 13, 2019

Abstract

This paper investigates how the transmission of monetary policy to the real economy depends on income and collateral borrowing constraints. Using an original dataset of UK mortgages, we assess the effects of monetary shocks on macroeconomic variables by exploiting time variation in the characteristics of new originations. We find that monetary policy is most potent when there is an abundance of mortgages recently issued at high income multiples. In contrast, when collateral constraints are loose and a large fraction of mortgages are issued at high loan-to-value ratios, monetary policy has a relatively stronger effect on asset prices. Taken together, we therefore provide new evidence for the importance of both the income and collateral channels of monetary policy.

Keywords: Monetary policy, credit conditions, mortgages, LTI, LTV, state-dependence.

JEL Codes: E52, E58.

1 Introduction

After almost a decade of near-zero policy interest rates, central banks around the world are once again starting to consider reactivating monetary policy. But the understanding of the monetary transmission mechanism that existed in the Great Moderation is now being called into question. A sustained period of low interest rates, high asset prices and higher income gearing (the level of borrowing against income) have the potential to change the effectiveness of monetary policy relative to that of the previous tightening cycles of the mid-2000s. This paper sheds light on how the transmission mechanism varies depending on the prevailing credit conditions in the economy.

It is by now well understood that the mortgage market is a key part of the transmission mechanism of monetary policy. As loan-to-value (LTV) constraints slacken, household budget constraints become more sensitive to asset price changes. This collateral channel of monetary policy can be thought of as the household equivalent of Tobin's Q and has been shown to amplify consumption responses to policy rate (e.g., Cloyne et al. (2016)). But the interaction of large debt contracts with changes in interest rates can also lead to important income effects that drive household behaviour. High loan-to-income (LTI) ratios are associated with household disposable incomes becoming more sensitive to changes in debt servicing costs. Despite many years of research and policy experience, these competing channels of monetary propagation are still difficult to disentangle.

This paper investigates to what extent the transmission of monetary policy to the real economy depends on income and collateral borrowing constraints. In particular, we take the distribution of LTIs and LTVs above given thresholds on the flow of new mortgages to proxy for the economy being more or less exposed to changes in interest rates. This is motivated by evidence that household behaviour might react in non-linear ways at very high levels of debt relative to income or collateral. By identifying periods where the real economy is more or less exposed to the two channels, we can assess the importance of each for the overall traction of monetary policy.

To understand the evolution of credit conditions, we exploit proprietary data on the monthly flow of the universe of UK regulated mortgages since 2005. This allows us to track the evolution of income gearing and collateral leverage with precision over time. By combining macroeconomic data with these novel loan-level mortgage data, we are able to estimate the state dependency of monetary policy at a granular level, and make progress in parsing out the contribution of collateral and income effects in the monetary transmission mechanism.

This paper uses UK data and standard identification techniques to form a set of monetary shocks. We then use local projections to investigate the importance of the flow of credit to households in the transmission of monetary policy shocks to the rest of the economy. Such local projections have been shown to be robust to misspecification and are a flexible way to generate multivariate impulse responses. In particular, we look at the response of real-economy variables to monetary policy shocks after a period of loose and tight credit conditions, defined by the flow of mortgages at various LTI and LTV thresholds.

We find that monetary policy has a more potent effect when a large share of households are close to their income borrowing constraints. This effect works through at least two channels. First, when households increase their income gearing, the mechanical effect of monetary policy on disposable income is amplified. Second, households closer to their borrowing constraints are likely to exhibit a high marginal propensity to consume. Taken together, we find that monetary policy in an economy closer to its income borrowing constraint has a larger effect on durable consumption and economic activity. In contrast, monetary policy has a larger effect on the economy through asset prices when households are closer to their collateral constraint.

One concern with our approach is that we may be capturing factors other than the state contingency of monetary policy to borrowing constraints. For example, in the last two decades, the UK has undergone significant structural changes along several dimensions including the composition of mortgages. To allay these concerns in our central specification we control 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 to the share of floating-rate mortgages in issuance. We also control for the benchmark reset interest rate and the overall flow of mortgage originations.

Another concern with our definition of high and low thresholds of LTIs and LTVs is that we might inadvertently be capturing the effect of 2008-2009 global financial crisis. Several things happened during that period, such as policy rates reaching the zero lower bound (ZLB) and a dislocation in the UK mortgage market (e.g. bank appetite to issue high LTV mortgages became severely impaired). We therefore run a placebo test and assess the time-dependence of monetary policy pre and post 2008-2009 and show that we can still successfully strip out the effects of collateral and income gearing even though the economy was undergoing important structural changes.

Our main finding can be illustrated with some context about how interest rates affected household finance during the Great Recession. When the Monetary Policy Committee reduced interest rates by over 400bp in 2008 the average UK mortgagor received a favourable cash-flow shock of around 5% of their annual pre-tax income. For those on high-LTI mortgages, the reduction in mortgage payments was closer to 8%, close to a quarter of their annual discretionary income. Monetary policy mechanically had its most direct impact on those households who had borrowed the most relative to their incomes. At the same time, house prices fell by several percentage points between the summer of 2008 and the spring of 2009. Those with recently-issued high LTV mortgages were the most at risk of being unable to refinance their mortgages as they approached negative equity. This means that they could have been trapped in mortgages with unfavourable terms until house prices recovered. These examples show that income and collateral leverage are likely to have non-linear effects on household's ability to smooth their consumption in the face of shocks. Although we exploit more moderate changes in conditions and behaviour for our analysis, the Great Recession paints a useful picture for why these constraints might matter.

Our paper adds to a rich literature that has established the important link between household balance sheets and the the propagation of shocks, both in terms of assets (e.g., Attanasio et al. (2002); Iacoviello (2005); Eggertsson and Krugman (2012); Mian and Sufi (2011); and Ragot (2014)) and in terms of income flows (e.g., Kaplan and Violante (2014); Auclert (2017); and Greenwald (2018)). Our paper builds on the recent works of Cloyne et al. (2016) to dig deeper into the overall effects of monetary policy and separate out the various channels that affect household behaviour.

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)). The UK mortgage market is well known to have short fixation periods and interesting dynamics on a range of metrics. In this study we therefore exploit the unique structure and data access of the UK mortgage market as a way in to understanding the precise 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), Aikman et al. (2017), Alpanda and Zubairy (2019), Ottonello and Winberry (2018), Beraja et al. (2017) and Cloyne et al. (2016)) 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 the 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. In relation to this literature, Tillmann (2017) examine the state-dependence of monetary policy to the level of uncertainty in the economy while Hubert (2018) investigates the effect of the central bank information channel on the monetary policy transmission.

The rest of this paper is organised as follows. In Section 2 we set out the two channels of monetary policy we aim to capture and describe the data that helps us construct our state variables. In Section 3 we go on to detail the empirical strategy we employ by first describing how we define monetary shocks and then setting out our local projection framework. We outline the main results and various robustness tests in Section 4, and Section 5 concludes.

2 Measuring Borrowing Constraints

2.1 The income and collateral channels

In this section we explore two alternative channels for the transmission mechanism of monetary that we attempt to isolate in this paper. Although monetary policy affects households along various dimensions, here we focus on the transmission of monetary policy through household debt, which is of primary importance to the UK.

The income channel of monetary policy captures the reallocation and general equilibrium effects of changes in interest rates on household resources and behaviour. A reduction in interest rates leads to a redistribution of net income from savers to borrowers, the latter of which are perhaps more likely to have higher marginal propensities to consume. Beyond direct cash-flow effects, the aggregate impact of lower interest rates is likely to affect borrowers and savers via boosts in other real income flows associated with higher economic activity (e.g., see Auclert (2017)).

Household income gearing is likely to refract the impact of changes to interest rates because it alters the proportionate effect on disposable income. Those spending a large fraction of their post-tax income on mortgage repayments will therefore be more sensitive to the direct and indirect effects of monetary policy. Since behaviour with respect to proportionate changes in disposable income is likely to be non-linear, a high fraction of mortgagors on high-LTI mortgages is likely to represent a household sector that is significantly more responsive to monetary policy shocks.

In contrast, the collateral channel captures the effects of monetary policy on the asset side of household balance sheets. Household collateral leverage refracts the change in interest rates because of the link between asset prices, household net worth and their capacity to borrow in the future, including to fund consumption.

This is compounded by a further income channel in the UK because there is a relationship between mortgage interest rates and LTVs (and remortgages are frequent). Higher asset prices therefore directly affect the cash-flows associated with mortgage repayments in the future. Reflecting the underlying risk characteristics to lenders, the link between interest rates and LTVs is especially convex at high levels of leverage. As a result, borrowers on high-LTV mortgages are likely to respond the most to changes in monetary policy.

2.2 Using PSD data to measure constraints

The mortgage data we use 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 the borrower, loan and house characteristics associated with

the 11 million residential mortgages issued between 2005 and 2017. Of those, around half 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 flow of different types of mortgages during the period that we study. In particular, we track the real-time proportion of mortgages issued at LTI ratios of more than 4 and LTV ratios of more than 90% every month between April 2005 and December 2017. Across the sample, around 15% and 8% of mortgages were issues with an LTI above 4 and an LTV above 90%, respectively. But these proportions varied over time with credit conditions and developments in asset prices. We exploit this variation and treat these two baseline metrics as proxies for fluctuations in credit conditions related to income gearing and collateral constraints.

Figure 1: LTI and LTV Distributions 2005-2017

(a) LTI Distribution   (b) LTV Distribution
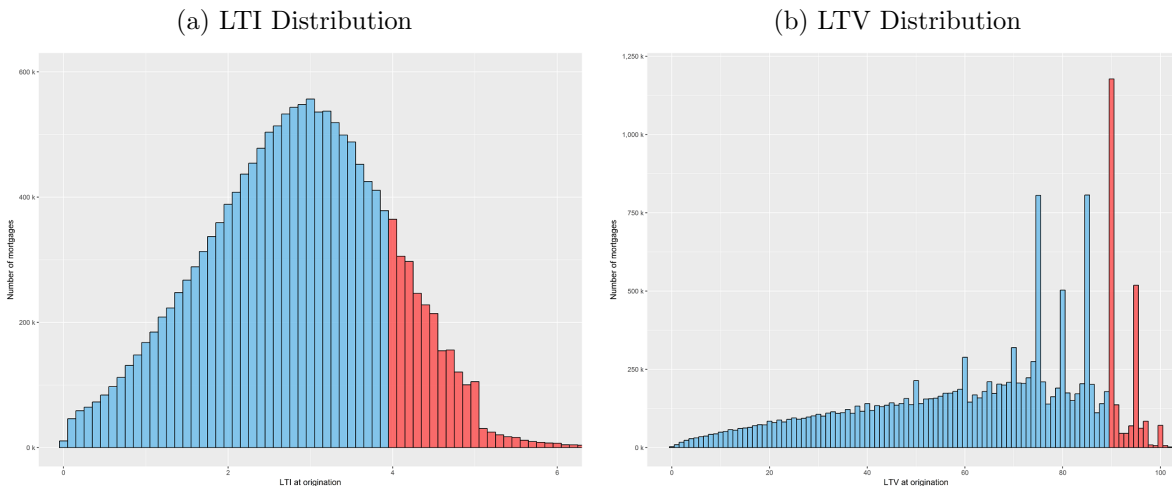
Note: Red bars correspond to the LTI cut-off of 4 and the LTV cut-off of 90%.

Figure 1 gives a sense of the the number of mortgages issued at different LTVs and LTIs across the entire sample. The chart on the left shows the distribution of LTIs is approximately normal and there is comparatively little bunching at various thresholds. The chart on the right shows there is much more bunching for LTV thresholds, especially at 5pp increments. Over our period, the modal mortgage was issued at an LTV of 90%.

The cut-offs we use for the proportions of high LTV and LTI mortgages are driven by an effort to capture a non-linear change in behaviour that 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. We therefore judge that the proportion of mortgages above an LTI of 4 does a reasonable job at capturing the non-linear sensitivity of households as a whole to changes in interest rates.

High LTV mortgages are often associated with higher interest rates in the UK to reflect the likelihood of the value of the collateral being insufficient to repay the liability in the event of a default. In 2017, the average interest rate for a (below) 90% LTV mortgage was around 2.8%. The equivalent rate for a 95% and and 85% LTV mortgage was 2.3 and 3.6%, respectively. At the other extreme, the interest rate gap between an new mortgage issued at 60 or 70% LTV was a mere 2bp. We therefore choose the 90% LTV mortgage as the threshold to capture the kink between interest rates and the stock, and net flow, of household resources, such that the assumption that mortgages above that threshold are constrained appears reasonable.

Figure 2: LTI State Variable
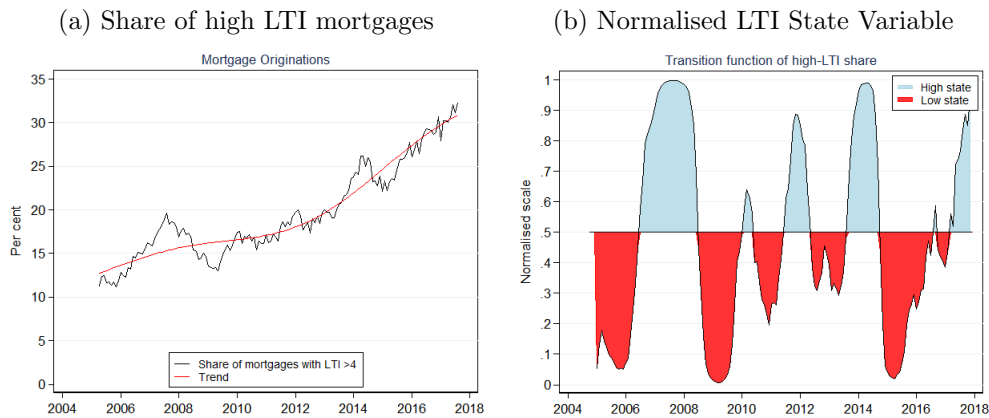
(a) Share of high LTI mortgages   (b) Normalised LTI State Variable

Note: Red shaded areas on the right panel correspond to regions where the blue line is below the red trend on the left panel. An analogous figure for the LTV state variable is shown in the Appendix.

Based on these threshold values for LTIs and LTVs, the evolution of the monthly proportion of new mortgages above them should reflect the degree to which borrowing constraints bind for the economy as a whole. In order to measure the effects of monetary policy as a function of the state of these constraints when the shock hits, we define $S(z_{i,t})$ as an indicator of the regime. $S(z_{i,t})$ is a smooth transition function of an indicator of the state of borrowing constraints $z_{i,t}$. This continuous variable captures the probability of the economy being in one of the two regimes: a high share of binding constraints or a low share. $z_{i,t}$ is the de-trended time series of the share of new mortgages above certain thresholds, changes in trend being a proxy measure of changes in the structure of the market or in preferences. We compute the following function:

$$S(z_{i,t}) = 1 - \frac{\exp(-\gamma z_{i,t})}{(1 + \exp(-\gamma z_{i,t}))}, \gamma > 0 \tag{1}$$

The variable $S(z_{i,t})$ is defined over the interval $[0 \ ; \ 1]$. We then refer to the period as being in the high state if the proportion of mortgages is above 0.5. This specification closely

resembles the smooth transition function of Granger and Terasvirta (1993). The parameter γ determines how abruptly the economy switches from one regime to the other when $z_{i,t}$ varies. This indicator of the state of borrowing constraints, $S(z_{i,t})$, is shown in Figure 2

The measures we construct use the proportion of new mortgages above these thresholds as a proxy for the sensitivity of the macroeconomy to interest rates. One concern might be that the stock of mortgages is large relative to the flow and we are therefore missing a significant part of the transmission mechanism in our analysis. Aside from data availability, there are good reasons to use the flow rather than the stock. Mortgagors tend to be most constrained early on in the contract. First time buyers are much more likely to have high-LTV mortgages and the flow of remortgages gives a good indication of how these new immature mortgages are evolving. For more seasoned mortgages, the LTVs are likely to be lower and increases in income associated with the life-cycle make mortgage repayments easier to bear. For these reasons, focusing on the flow of mortgages gives us a cleaner read on the most constrained parts of the household sector.

2.3 Discussion

Our choice of using thresholds of LTIs and LTV is our best attempt at identifying the state contingency of the effect of monetary policy. As discussed above, both of these metric capture slightly different channels of transmission. The LTI metric captures the income channel, which can loosely be thought of a mixture of cash-flow and general equilibrium effects. The LTV metric captures a collateral channel and a milder form of the income channel in terms of the cash-flows associated with mortgage payments in the future.

For this reason there might be some concern that the two metrics are picking up the same variation and are therefore correlation. In the robustness checks we make sure the results hold through when we split the sample into four buckets with the permutations of high and low for each measure. But we can also use the microdata to show that the two measures are first-order independent. Table 1 shows a loan-level regression explaining LTV and LTI using a large subset of variables available in the PSD. While the coefficients are interesting on their own, the focus of this table is the last two rows, which show that the correlations between the two variables, and residuals of the two specifications, are both very close to zero.

An alternative metric to the LTI metric of a mortgage is the payment-to-income (PTI) ratio. PTIs reflect the ability of households to service debt payments and are therefore closely associated with LTIs. But the evolution and interpretation of PTIs has two complicating factors between 2005 and 2017, which make them unattractive to use as a state variable. First, interest rates have steadily fallen over the majority of the post-crisis sample, which mechanically pushes down on trend PTI. Second, mortgage terms have steadily increased as households have attempted to afford higher house prices even in the face of lower mortgage rates. Since UK institutions and regulators monitor the flow of LTI mortgages and put less weight on PTIs, we argue that the evolution of LTIs is more informative to cleanly identify our channels of interest. Figure 10 in the appendix demonstrates that the PTI distributions

are very different to the LTI distributions in 2007 and 2017.

Table 1: The Independence of Borrowing Constraints

| | Loan-level regressions | |
	LTV	LTI
Property value (£000s)	-0.41***	0.035***
	(0.002)	(0.000)
Income (£000s)	1.81***	-0.021***
	(0.002)	(0.000)
Age	-0.80***	-0.03***
	(0.001)	(0.000)
Mortgage type (7 categories)	Yes	Yes
Interest rate type (6 categories)	Yes	Yes
Region (12 categories)	Yes	Yes
Time dummies (54 categories)	Yes	Yes
Mortgage type (7 categories)	Yes	Yes
Observations	16,796,188	16,796,188
R-squared	0.12	0.06
Raw data correlation	-0.03	
Cross-residual correlation	0.14	

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

Finally, it is worth exploring some additional issues that we have addressed in constructing and exploring our state variables. The regression in Table 1 provides some evidence that the two metrics are not capturing the same information. Nevertheless, in the robustness tests we include four-way bins that interact the high and low states for the LTI and LTV thresholds. These results help us separate out the different channels of monetary policy. But there might be some remaining concerns that groups or regions of unrepresented households may dominate some of the effects we find. We therefore run specifications where we drop London, which is known to have experienced idiosyncratic conditions to the rest of the UK after 2009. To allay fears that some of our time periods are unrepresentative, we construct a version of the state variable trend with a structural break at the start of the Great Recession. These complementary results are explored later on the paper.

3 Empirical Strategy

3.1 The identification of monetary shocks

Because we aim to investigate the causal impact of monetary policy on various financial and macroeconomic variables, we face the usual endogeneity issue of the policy instrument. More specifically, the monetary policy instrument is set in response to developments in financial and macroeconomic variables, but in the meantime, these variables respond to changes in

the policy instrument. There is an abundant literature focusing on different ways to solve this reserve causation problem. Three strands have emerged: the use of timing restriction and a Cholesky decomposition of the matrix of the residuals from a structural VAR, high-frequency event-study strategies, or the Romer and Romer (2004) approach that consists in singling out the systematic endogenous part of monetary policy.

The first approach has been shown to produce the so-called puzzle, and if the econometrician's information set is expanded, to reduce the number of degrees of freedom. The second approach works well with high-frequency data, but the mortgage data used for the identifying borrowing constraints are available at the monthly frequency. The third approach mixes three advantages: it is appropriate for low-frequency data and relatively short samples, it is quite parsimonious in the number of parameters to estimate, and it enables controlling specifically for the information set available to policymakers at the time of the decision. Romer and Romer (2004), Coibion (2012) and Cloyne and Huertgen (2016) have shown that this identification strategy works well to investigate the transmission of monetary policy to the real economy.

The Romer and Romer (2004) has first been applied to UK data by Cloyne and Huertgen (2016). 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:

$$\Delta i_t = f(\Omega_t) + \epsilon_t^i \tag{2}$$

Changes in the policy instrument, Δi_t are regressed on the central bank information set Ω_t at the time of the decision. The residual, ϵ_t, represents the monetary shock series. Blanchard et al. (2013) and Miranda-Agrippino and Ricco (2017) have shown 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 ?? 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 standard policy rate, Bank rate in the UK, may not be relevant to measure the stance of monetary policy over the sample period considered. First, because of the Zero Lower Bound (ZLB) and the fact that monetary policy has become multi-dimensional since the Great Recession with the implementation of asset purchases, forward guidance or targeted liquidity provisions, the policy rate is not relevant enough to capture the overall stance of monetary policy. Second, a large consensus has formed about the content of monetary policy news: the main piece of information on central bank announcement days relates to changes in the future likely policy path, as suggested by Gürkaynak et al. (2005), Campbell et al. (2012) and Hanson and Stein (2015). The primary share of the news contained in policy announcements is about the expected path of future policy (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 deal with 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. We follow Hanson and Stein (2015) and use 2-year interest rates. The key point is that this measure could capture revisions in private beliefs about the expected medium-term policy path as opposed to contemporaneous policy decision only, and to encompass in one single variable conventional and unconventional policies.

Equation 2 can thus be rewritten as follows:

$$\Delta i_t = \alpha + \rho i_{t-1} + \sum_{j=1}^{3} \lambda_j \mathbb{E}_t^{CB} \pi_{t+j} + \sum_{j=1}^{3} \gamma_j \mathbb{E}_t^{CB} x_{t+j}$$
$$+ \sum_{j=1}^{3} \phi_j \Delta \mathbb{E}_t^{CB} \pi_{t+j} + \sum_{j=1}^{3} \varphi_j \Delta \mathbb{E}_t^{CB} x_{t+j} \qquad (3)$$
$$+ \sum_{k=0}^{3} \psi_k U_{t-k} + \kappa \mathbb{I}_t^{IR} + \eta_{swap} \mathbb{E}_t^{swap} \pi_{t+2} + \eta_{sef} \mathbb{E}_t^{sef} \pi_{t+2} + \epsilon_t^i$$

where i_t is the policy instrument, measured with 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, respectively, at horizons 1, 2 and 3 years ahead, U_{t-k} is the unemployment rate for different months k, $\mathbb{I}_t^{IR}$ is a dummy that takes the value 1 in months when the Inflation Report is published[1], $\mathbb{E}_t^{swap} \pi_{t+2}$ is a measure of inflation expectations at the 2-year horizon from financial markets 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.[2] 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. The residuals ϵ_t^i are the proxy for the monetary shock series. Figure 11 plots this estimated monetary shock series. As expected, the largest values happen around 2008 and 2009 when the Great Financial Crisis was hitting the UK economy. We have also tested that these monetary shocks are unpredictable from movements in macroeconomic data.[3] In the robustness section, we verify that our main results are not sensitive to our proxy for the policy instrument, using a shadow rate, or to the ZLB period.

---

[1]One may argue that policymakers is more likely to update policy when they update their macroeconomic projections, so in Inflation Report months. This is not the case over this sample, but we control for this potential effect.

[2]We consider two different types of inflation expectations, from financial 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.

[3]These tests are available from the authors upon request.

Figure 3: Monetary policy shocks
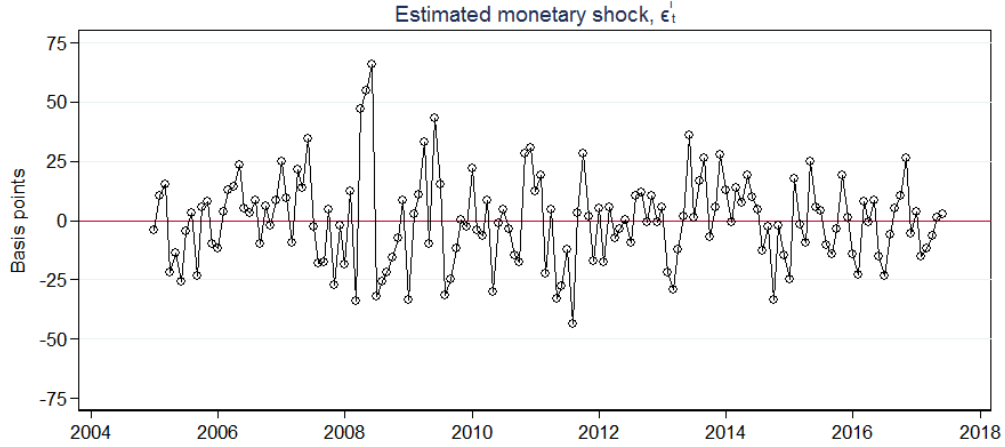
Note: Our preferred monetary shock series, estimated with equation 3, refers to the change in the two-year risk-free interest rate.

3.2 Local projections

We aim to measure the dynamic effect of monetary shocks conditional on the share of new mortgages with LTI (or LTV) above a certain threshold. To do so, our preferred approach is to use local projections as proposed by Jordà (2005) with our externally identified instruments for monetary shocks. We interact the monetary shock series with the state-variable S_t described in the previous section. This Jordà (2005) method has become a very popular tool to compute impulse responses because of its robustness 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 can easily account for non-linearities in the transmission of monetary policy. 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 autoregressive distributed lag model. Equation 4 is therefore estimated k times as follows:

$$y_{t+k} = \alpha_k + \beta_k \epsilon_t^i + \gamma_k S_t + \delta_k \epsilon_t^i \cdot S_t + \phi_k X_t + \epsilon_t \tag{4}$$

where ϵ_t^i is the monetary shock series, S_t the state-variable capturing borrowing constraints, and $\epsilon_t^i \cdot S_t$ their interaction. X_t is a vector that includes three lags of the dependent variable y_t, of the policy rate i_t and of the variable measuring the state of the economy, namely the unemployment rate. X_t also includes the average mortgage duration and the share of floating mortgages in order to control for changes in two crucial features of the structure of the mortgage market over our sample. Finally, because credit conditions are endogenous to the monetary transmission mechanism, changes in policy rates gain traction on the economy by influencing the price and quantity of loans available to households and firms (e.g. the External Finance Premium). Thus, X_t also includes the standard variable

11

rate, the rate at which mortgages revert to absent a refinance, and the total flow of new mortgages originated at each period in order to control for price and quantity dynamics of the mortgage market. Equation 4 is estimated using OLS over the sample going from April 2005 to December 2017. We compute heteroskedasticity robust standard errors.

4   The Transmission of Monetary Policy

In this section of the paper we set out our baseline results that show the different channels of the monetary transmission mechanism that can be captured using data on credit conditions at different point in the business cycle. In a standard treatment of policy transmission, aggregate credit conditions have little effect on the behaviour of the economy. In the following analysis we would therefore expect to see that monetary shocks have the same effect on all financial and real-economy variables regardless of the prevailing credit conditions. In other words, we should expect to see overlapping trajectories for responses in the face of a 1pp shock in the policy rate.

4.1   The role of income constraints

We find that the effect of monetary policy on financial and macroeconomic variables depends on the flow of high-LTI mortgages. In Figure 4 the impulse responses of monetary policy shocks are represented in the beige swathe in the presence of a relatively large number of high-LTI mortgages using our definitions above. In contrast, the blue swathes show the effect of monetary policy when credit conditions are tighter, and the proportion of high-LTI mortgages lower.

As shown in the first row of Figure 4, a one percentage point monetary shock policy shock, conditional on recent pickup borrowing relative to income, leads to a peak impact on dampened financial market inflation expectations of around 20bp after four quarters. The effect on realised inflation measures is stronger still, as the peak impact on producer and consumer prices is around 4pp and 1.8pp, respectively. Monetary policy also has a state contingent effect on real-economy variables in the presence of high-LTI mortgage issuance. Durable consumption, industrial production and employment all intuitively decline after a monetary policy tightening. Impulse responses for other variable are shown in the Appendix.

The gaps between the beige and blue swathes provide evidence that the effectiveness of monetary policy might vary according to how close an economy is to its borrowing constraints. When an economy has a high proportion of new mortgages issued at high multiples of income more households are likely to be closer to the limits of what they can afford to service. These households, and the economy as whole, should be expected to be more sensitive to changes in interest rates. We therefore use the LTI threshold to proxy for the income channel of monetary policy. An example of one direct channel is through spending decisions after a change in mortgage repayments, though the graphs suggest there might be broader

effects over and above the pure liquidity or cash-flow channel.

The blue swathes in Figure 4 show counter-intuitive responses of several variables following a monetary policy shock in the presence of a low issuance of high-LTI mortgages. This might be surprising because although high income gearing might magnify responses, at first glance it is not obvious why smaller income shocks reverse the direction of agents' responses. We argue there are at least two plausible explanations for the 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 a one percentage point shock 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. For example, if a monetary shock reveals the policymaker's view on the outlook for the economy agents might respond in ways consistent with the blue swathes.

Figure 4: Monetary Transmission in the Presence of High-LTI Mortgages
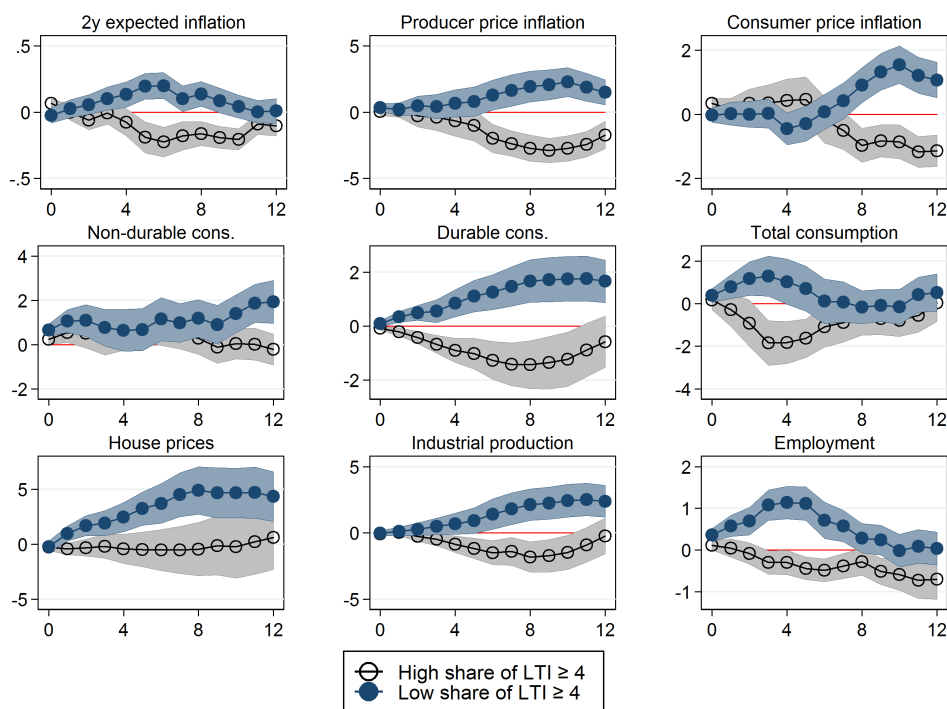
Note: This figure shows the estimates of the effect of ϵ_t^i over 12 months for different macroeconomic variables, based on the OLS estimation of equation 4 over the sample April 2005 - December 2017. The black empty circle corresponds to the effect of a contractionary monetary policy with a high share of LTI above 4 whereas the blue full circle corresponds to the effect of a contractionary monetary policy with a low share of LTI above 4. The shaded areas represent 1 SE.

4.2 The role of collateral constraints

We find slightly less evidence that the transmission of monetary policy is varies as much through the collateral channel. Figure 5 shows the impulse responses when we use the flow of high-LTV mortgages as the state variable in the local projections. Once again the inflation variables respond in the usual directions in the beige swathes, with similar profiles to those in Figure 4. The broad paths also show approximately the same shape for the macroeconomic variables, though the gap between the beige and blue swathes is, on balance, slightly smaller.

Of the nine panels in Figure 5, the most strikingly different to Figure 4 is the response of housing. When high-LTV mortgages are used as the state variable, the transmission of monetary policy through to house prices become more state contingent. In particular, the beige swathe stays close to or below zero for the first eight quarters before picking up at the end. At the same time the blue swathe climbs rapidly over the first eight quarter before falling back. This apparent state contingency of the response of house prices to monetary policy is consistent with the collateral channel of monetary policy. A large flow of recently issued highly-levered mortgages is likely to be a proxy for an economy bumping up against a collateral constraint, which makes behaviour more sensitive to changes in financing conditions. On the other hand if there is collateral-leverage slack in the economy interest rates are likely to have less potency in depressing asset price bubbles, especially when those assets are purchased using leverage.

Overall, our results using high-LTV mortgages provide evidence that monetary policy is more effective in the presence of collateral constraints, but that it is perhaps relatively more effective in its effect on real-economy asset prices. Taken together, Figures 4 and 5 provide evidence for both the income and collateral channel of monetary policy.

4.3 Monetary policy before and after the Great Recession: A placebo test

In our central specification we use the presence of high-LTI and high-LTV mortgages to proxy for income and collateral constraints that might lead to changes in 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 variables are 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.

Figure 6 shows the impulse responses when we split the sample into these two periods. The beige 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.

Figure 5: Monetary Transmission in the Presence of High-LTV Mortgages
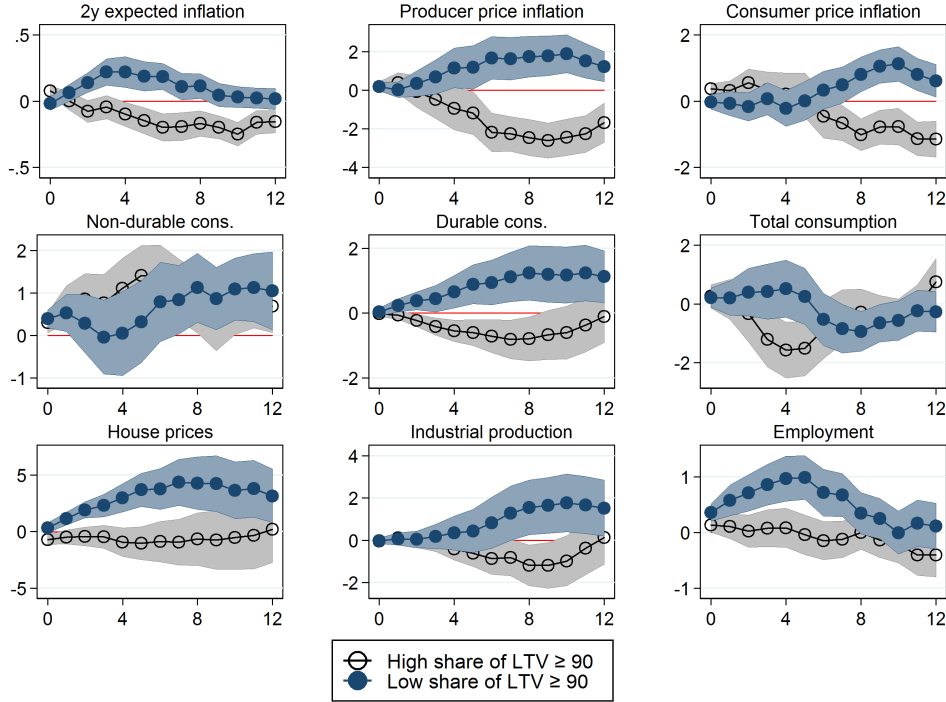
Note: This figure shows the estimates of the effect of ϵ_t^i over 12 months for different macroeconomic variables, based on the OLS estimation of equation 4 over the sample April 2005 - December 2017. The black empty circle corresponds to the effect of a contractionary monetary policy with a high share of LTI above 4 whereas the blue full circle corresponds to the effect of a contractionary monetary policy with a low share of LTI above 4. The shaded areas represent 1 SE.

The main takeaway from Figure 6 is that monetary policy has had slightly more traction in the period after 2008, but this cannot explain the profiles we observe in Figures 4 and 5. Although the response of inflation is stronger after 2008, the blue inflation swathes still dip below the zero line. The consumption responses are also different, with the placebo showing a significant change in non-durable consumption rather than durable consumption. Finally, the profiles in the last row do not support the hypothesis that the only effect we are capturing in the main specifications is purely to do with timing.

4.4   Extensions

Our results are robust to a number of variations around our main specification. Our central results rely on well-identified monetary shocks being used to disentangle the effect of monetary policy on various economic variables. Although there is good justification to use the surprise element of the two-year risk-free 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

Figure 6: Monetary Transmission Before and After the Great Recession
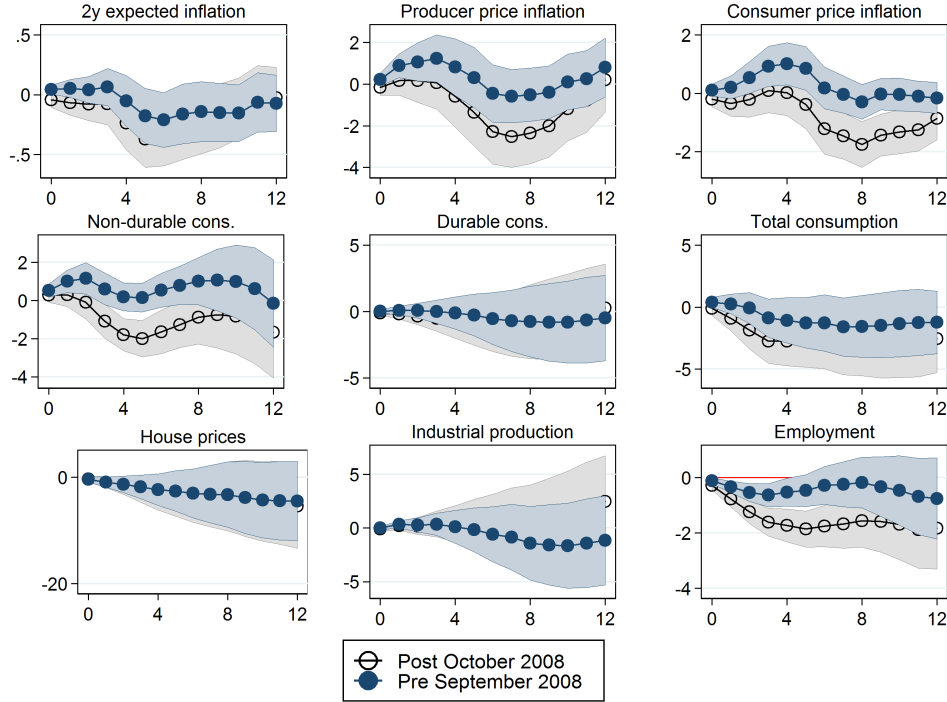
Note: This figure shows the estimates of the effect of ϵ_t^i over 12 months for different macroeconomic variables, based on the OLS estimation of equation 4 over the sample April 2005 - December 2017. The black empty circle corresponds to the effect of a contractionary monetary policy with a high share of LTI above 4 whereas the blue full circle corresponds to the effect of a contractionary monetary policy with a low share of LTI above 4. The shaded areas represent 1 SE.

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 12 shows remarkably consistent results. In both cases the directions of variable responses go in the right direction even if the state contingency is not always statistically significant.

The same broad pattern of results also holds when we vary the conditioning assumptions used in the construction of the monetary shock. For the overall directions of the responses, it seems hardly to matter which financial market prices we use, or which forecasts we condition on. In more standard narrative approaches it is very important to condition on the appropriate information set of the central bank. It turns out that various assumptions can be used to produce results in line with our hypothesis on the income and collateral channel of monetary policy.

To ensure that our main results are giving a plausible sense of the transmission of monetary policy, in Figures 13 we show the impulse responses of other variables that can be used

as a cross-check. The first rows show the initial transmission through to financial markets appears to be in the right direction and of the correct magnitude. The second rows show the response of the stock market and other inflation-measure variants that respond intuitively. Finally, the last rows show some additional real-economy variables that behave sensibly.

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 and LTV thresholds represent an economy close to the income and collateral constraints. This is relatively straightforward and the values of 4 and 90% make sense in terms of the context of mortgage market practices and regulation. All else equal, Figure 14 shows that we get similar results using slight variations on these thresholds.

One concern with the share of mortgages with a LTI above a certain threshold could be that it would capture households with strong preference for debt more than households close to their income constraints. For instance, households with very high revenues could have a very high LTI but still a comfortable disposable income after mortgage payments. We therefore compute the share of mortgages that interact the high and low states for the LTI and LTV thresholds, such that these households close to the two constraints are likely to be effectively close to their budget constraint. Another way to alleviate this potential concern is to compute the share of LTI above 4 while excluding all mortgages originated in the London area, which is known to concentrate most of the high-income high-wealth households and to experience idiosyncratic housing conditions to the rest of the UK. Figures 16 and 15 plot a similar state-dependent transmission of monetary policy tho the real economy.

The more subtle issue is to determine at which point there is a high or a low issuance of mortgages at these levels. Our baseline results examine the proportion of mortgage at these thresholds relative to the trend. Those results remain relatively constant whether we use a trend calculated across the whole sample or one only using information captured in real time. But our results also look qualitatively similar when we impose structural breaks in those trends, when we impose different assumptions about how smooth the trends are, and whether we choose our state variable to be continuous or discrete.

5 Conclusion

It is well-known that households responds more to the effects of monetary policy when they are closer to their borrowing constraints. Nevertheless, our findings provide some of the first empirical evidence for how changes in the interest rate feed through to different parts of the economy via these constraints. Moreover, we show that the transmission of monetary policy depends significantly on those constraints and the properties of the credit flows the economy is experiencing. Although monetary policy has long and variable lags, we show that credit conditions might explain a reasonable size of that variation.

Changes in monetary policy mechanically have a larger impact on disposable income

when households borrow more relative to their income. We argue that the additional increase in the marginal propensity to consume for those close to their budget constraints increases the potency of monetary policy on behaviour. We find that economic activity for an economy close to its income borrowing limit correspondingly responds more to monetary policy shocks. Intuitively, when an economy is closer to its collateral limit, the effect of monetary policy on asset prices, including housing, is higher. Together, this evidence could be used by policymakers to calibrate the response of monetary policy to various adverse shocks, as well as helping to coordinate monetary and macro-prudential policy in the face of multiple shocks and policy remits.

References

Aikman, D., Liang, N., and Modugno, M. (2017). Credit, financial conditions, and monetary policy transmission. Technical report, Hutchins Center Working Paper.

Alpanda, S. and Zubairy, S. (2019). Household debt overhang and transmission of monetary policy. *Journal of Money, Credit and Banking*.

Angrist, J., Jordà, Ò., and Kuersteiner, G. (2018). Semiparametric estimates of monetary policy effects: string theory revisited. *Journal of Business & 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. (2017). Monetary policy and the redistribution channel. NBER Working Paper 23451, NBER.

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. (2017). Regional heterogeneity and monetary policy. NBER Working Paper 23270, NBER.

Bernanke, B., Boivin, J., and Eliasz, P. (2005). Measuring the effects of monetary policy: A factor-augmented vector autoregressive (favar) approach. *Quarterly Journal of Economics*, 120(1):387–422.

Blanchard, O., L'Huillier, J.-P., and Lorenzoni, G. (2013). News, noise, and fluctuations: An empirical exploration. *American Economic Review*, 103(7):3045–70.

Caggiano, G., Castelnuovo, E., and Nodari, G. (2017). Uncertainty and monetary policy in good and bad times. *mimeo*.

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*, pages 1–80.

Cloyne, J., Ferreira, C., and Surico, P. (2016). Monetary policy when households have debt: new evidence on the transmission mechanism. Bank of England Working Paper 589, Bank of England.

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, Bank of England.

Di Maggio, M., Kermani, A., Keys, B. J., 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.

Finck, D., Schmidt, J., and Tillmann, P. (2018). Mortgage debt and time-varying monetary policy transmission. Magks papers on economics, Philipps-Universität Marburg, Faculty of Business Administration and Economics, Department of Economics (Volkswirtschaftliche Abteilung).

Flodén, M., Kilström, M., Sigurdsson, J., and Vestman, R. (2017). Household debt and monetary policy: Revealing the cash-flow channel.

Garcia, R. and Schaller, H. (2002). Are the effects of monetary policy asymmetric? *Economic Inquiry*, 40(1):102–119.

Granger, C. and Terasvirta, T. (1993). *Modelling non-linear economic relationships*. Oxford: Oxford University Press.

Greenwald, D. (2018). The mortgage credit channel of macroeconomic transmission.

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*.

Hanson, S. and Stein, J. (2015). Monetary policy and long-term real rates. *Journal of Financial Economics*, 115(3):429–448.

Hubert, P. (2018). State-dependent effects of monetary policy: the central bank information channel. *mimeo.*

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. and Violante, G. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

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. Technical report, OFCE Working Paper.

Ottonello, P. and Winberry, T. (2018). Financial heterogeneity and the investment channel of monetary policy. NBER Working Paper 24221, NBER.

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. Technical report, 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*, pages 85–108.

Appendix

Figure 7: LTI Minimum and Maximums

(a) LTI Min - April 2005
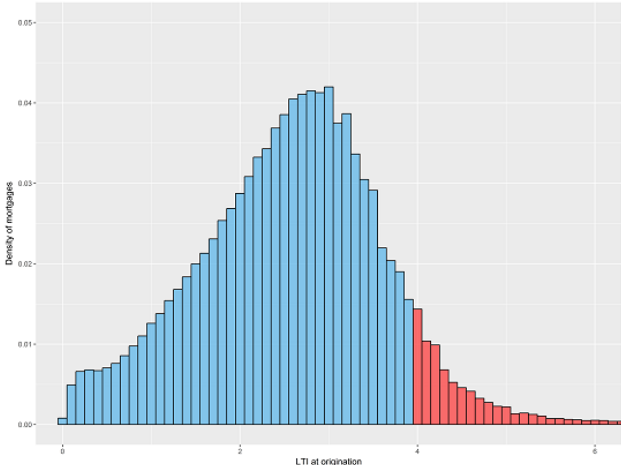
(b) LTI Max - August 2017
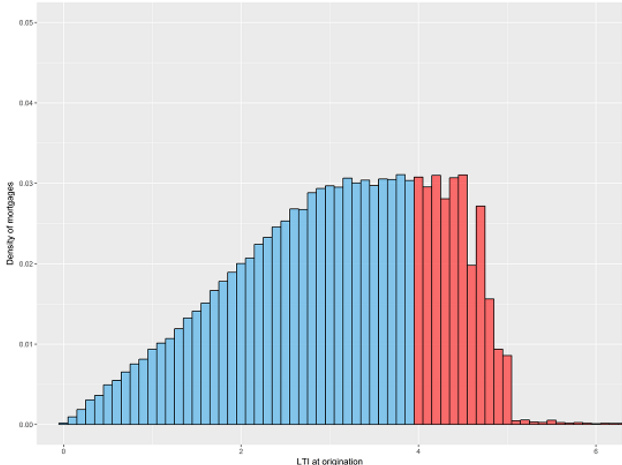
Note: Red bars correspond to cut-offs of 4 for LTIs.

Figure 8: LTV Minimum and Maximums
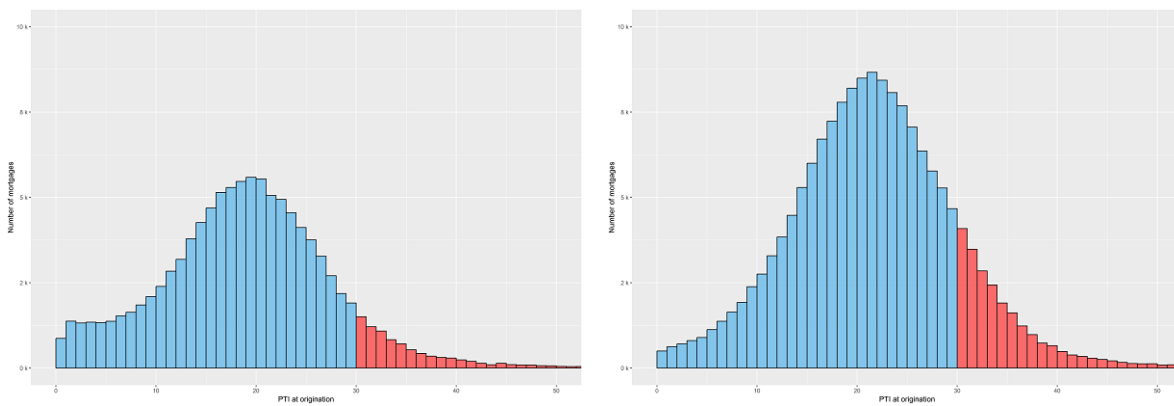
(a) LTV Min - March 2011
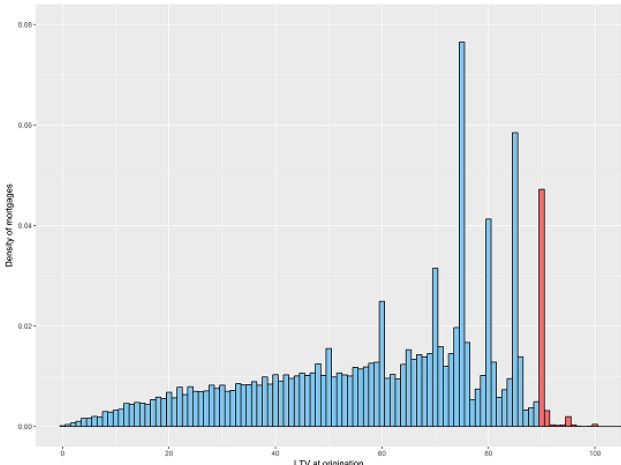
(b) LTV Max - June 2005
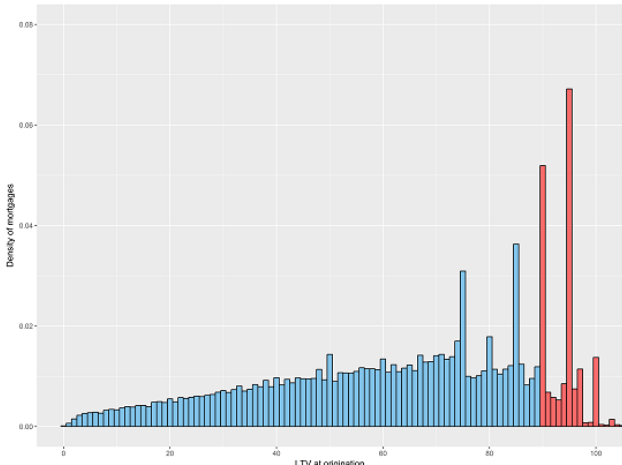
Note: Red bars correspond to cut-offs of 90% for LTVs.

Figure 9: LTV State Variable
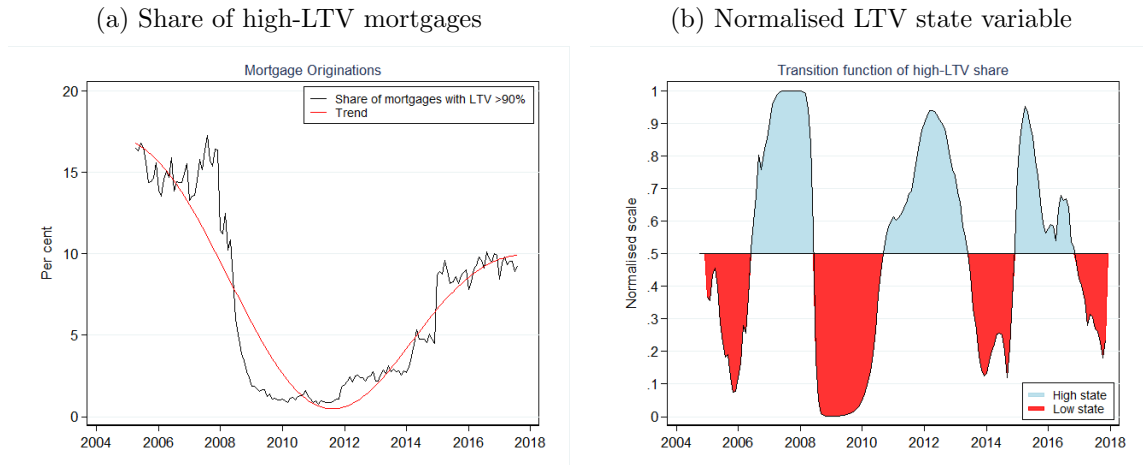
(a) Share of high-LTV mortgages

(b) Normalised LTV state variable

Note: Red shaded areas on the right panel correspond to regions where the blue line is below the red trend on the left panel.

Figure 10: PTI in 2005 and 2017

Note: Red bars correspond to cut-offs of 30% for PTI. The dates chosen correspond to those for the LTI minima and maxima in Figure 7

Figure 11: Monetary Policy Shocks
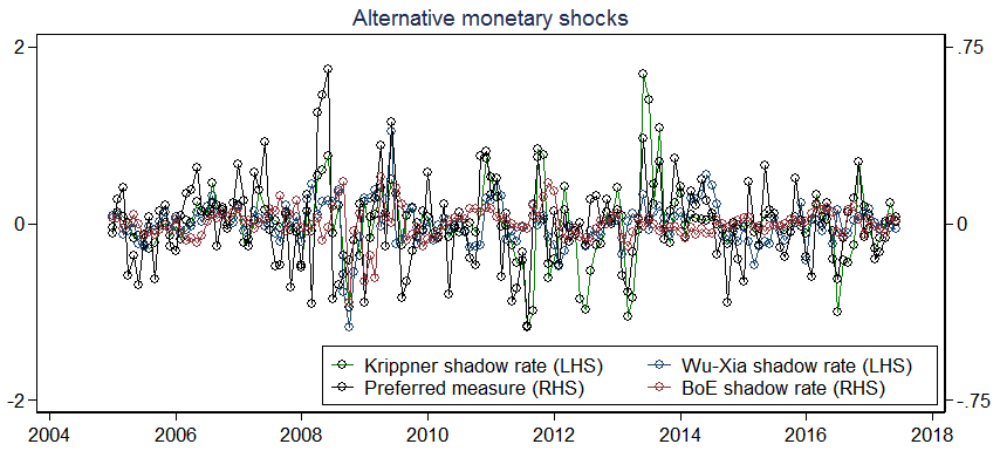
Note: This figure shows our preferred monetary shock in red.

Figure 12: Monetary Transmission with Alternative Shock Identification
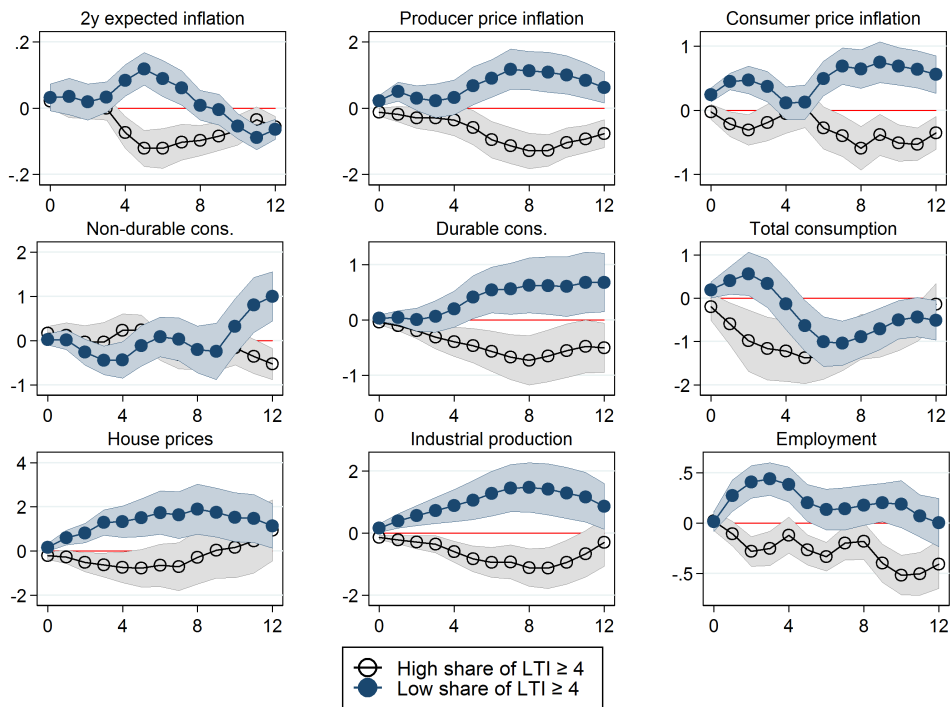
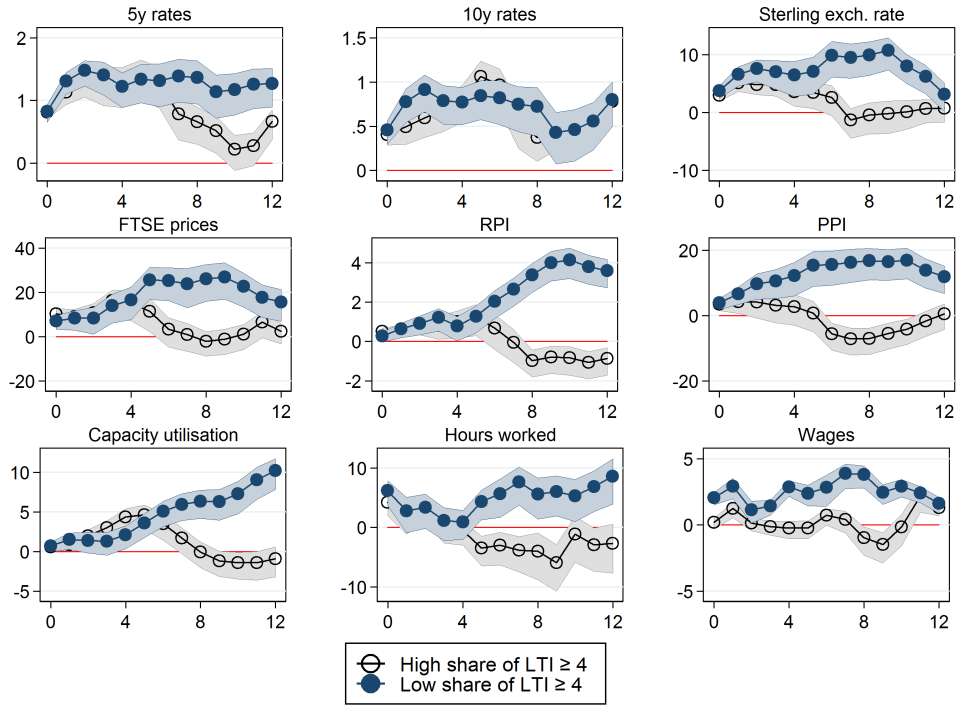Figure 13: Monetary Transmission for Other Variables
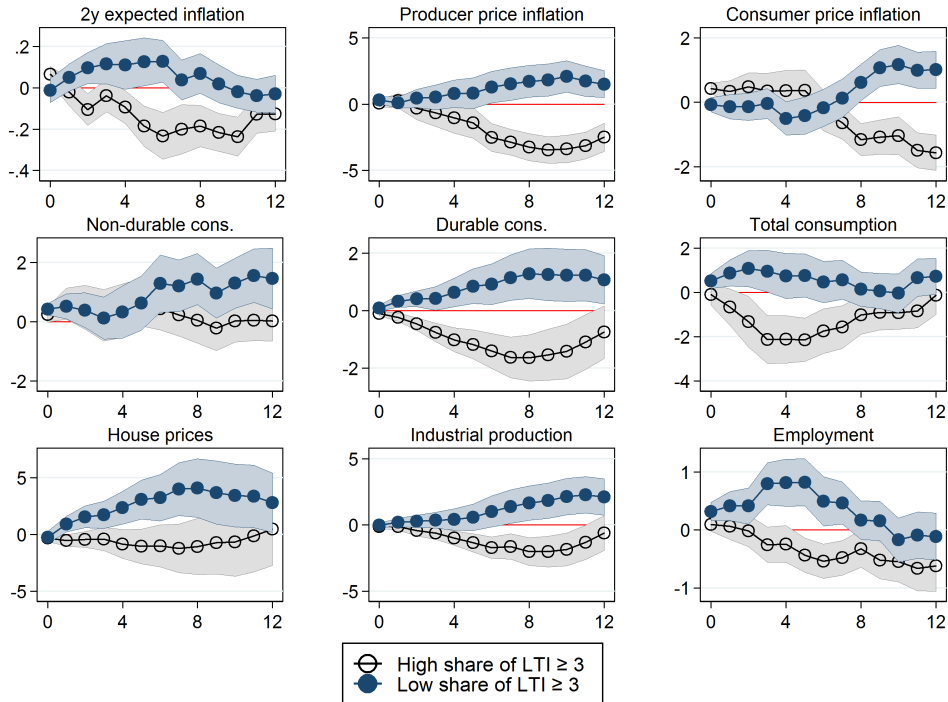

Figure 14: LTI state threshold of 3

Figure 15: Monetary transmission when excluding mortgages originated in London area
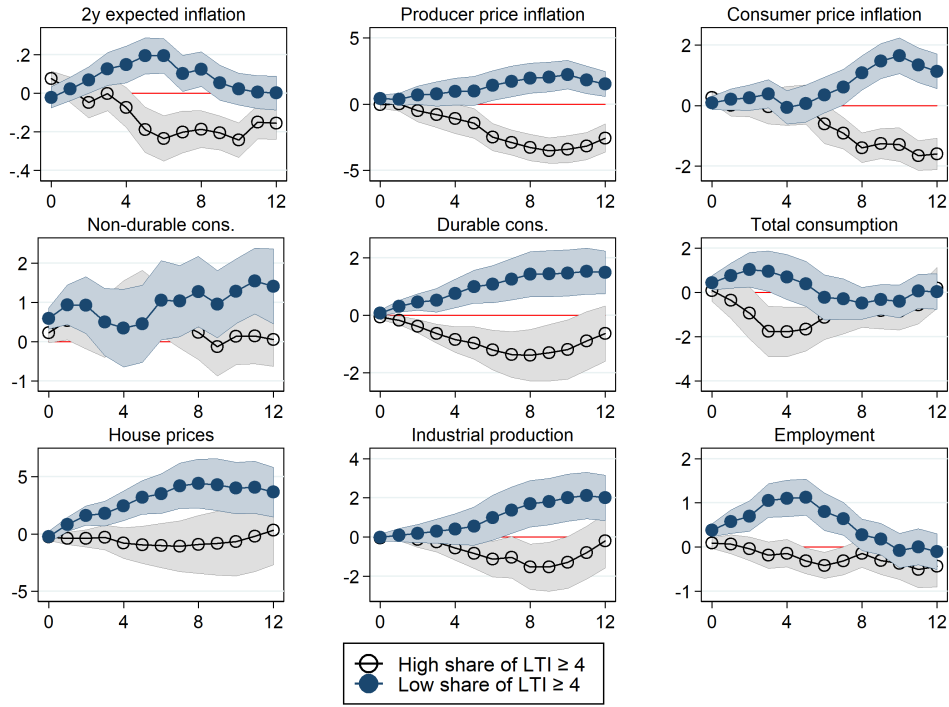

Figure 16: Monetary Transmission conditional on the share of high LTI and high LTV