Macroeconomic Dynamics (2023), **27**, pp. 506–531
doi:10.1017/S136510052100050X

ARTICLE

Mortgage debt and time-varying monetary policy transmission

David Finck*, Jörg Schmidt, and Peter Tillmann

University of Giessen, Germany.
*Corresponding author. Email: david.finck@wirtschaft.uni-giessen.de

Abstract
This paper studies the role of monetary policy for the dynamics of US mortgage debt, which accounts for the largest part of household debt. A time-varying parameter vector autoregressive (VAR) model allows us to study the variation in the sensitivity of mortgage debt to monetary policy. We find that an identically sized policy shock became less effective over time. We use a dynamic stochastic general equilibrium model to show that a fall in the share of adjustable rate mortgages (ARMs) can replicate this finding. Calibrating the model to the drop in the ARM share since the 1980s yields a decline in the sensitivity of housing debt to monetary policy which is quantitatively similar to the VAR results. A sacrifice ratio for mortgage debt reveals that a policy tightening directed toward reducing household debt became more expensive in terms of a loss in employment. Counterfactuals show that this result cannot be attributed to changes in monetary policy itself.

Keywords: Mortgage debt, monetary policy, debt reduction, time-varying VAR, DSGE

JEL Classfications: E3; E5; G2

1. Introduction

The build-up of household debt in the USA and other countries is often interpreted as a potential risk to financial stability (see Jordà et al., 2016) and a determinant of the overall credit cycle (see Mian et al., 2017). Since the recent financial crisis originated in the US housing market, the mortgage market receives a lot of attention. Mortgage debt is by far the largest component of household debt, as it reflects the single, most important financial decision of most households. Monetary policy should affect not only the value of houses, but also the dynamics of mortgage debt. A monetary tightening should curb the build-up of mortgage debt. This is one important channel for monetary policy transmission to households.[1]

In this paper, we study the time-varying sensitivity of US mortgage debt to monetary policy. To do so, a model that allows for time variation in the link between the Fed and the mortgage market is needed. We therefore use a time-varying parameter vector autoregressive (TVP-VAR) model along the lines of Primiceri (2005). The choice of this model is based on two observations: first, there is strong evidence that the US economy underwent both structural and institutional changes and also faces structural changes in the world economy, as emphasized by Canova and Gambetti (2009), Boivin (2006) and Mishkin (2009). Second, financial liberalization and deregulation changed the process of financial intermediation in the US economy. These structural changes cannot be picked up by a linear model with constant VAR parameters and a constant variance–covariance matrix. As a result, there is a danger of incorrect policy implications for central banks if the econometric model is not flexible enough to detect possible time-varying effects. We therefore rely on a model which allows for drifting coefficients and a time-varying variance–covariance matrix.

A model with discrete breaks would be an alternative estimation framework to describe the sensitivity of household mortgages to monetary policy shocks. However, a TVP-VAR is far more flexible than a model with discrete breaks and thus captures any form of structural shifts.[2]

Our model includes four variables: civilian unemployment, GDP deflator inflation and the deviation of real debt from its trend, intended to represent the non-policy block, and a short-term interest rate representing the monetary policy instrument. We assume that the Fed only responds to inflation and employment and restrict all time-varying VAR coefficients in the policy rule other than those related to inflation and employment to zero. Our estimation procedure relies on Markov chain Monte Carlo (MCMC) methods as in Nakajima (2011).

Our key result is that the reaction of mortgage debt to an identically sized monetary policy shock became much smaller over time.[3] Hence, a monetary policy tightening today reduces household debt much less than the same shock in the 1970s. For instance, a 25bp tightening in 1960 led to a drop in the cyclical component of mortgage debt by about 0.09 percentage points after eight quarters, while the same shock today would result in an insignificant reduction of mortgage debt. Importantly, we show that non-mortgage debt of households does not exhibit such a decline in its sensitivity to monetary policy.

We invoke the concept of the sacrifice ratio to shed light on when a debt reduction as a result of a monetary tightening is relatively expensive or cheap, respectively. We find that toward the end of the sample, it is particularly costly in terms of employment to use a monetary tightening in order to initiate a reduction in household debt.

To rule out that the change in the transmission to mortgages stems from shifts in the systematic behavior of monetary policy itself, we simulate counterfactuals in the spirit of Sims and Zha (2006) and keep the reaction function of the central bank fixed over time. The results are indistinguishable from our baseline finding. Hence, our results fit into the "mortgage rate conundrum" put forward by Justiniano et al. (2017). They argue that the link between Treasury yields, which party reflect monetary policy, and rates on mortgages weakened over time. Hence, both papers stress the role of some underlying structural changes in the mortgage market that impair the monetary policy transmission mechanism.

In order to understand our finding, we study the interaction of the interest rate elasticity of mortgages with the share of adjustable mortgages (ARMs). The literature points to the ARM share as a key determinant of policy transmission to the housing market (see, among others, Calza et al., 2013; Ehrmann and Ziegelmeyer, 2017; Garriga et al., 2017). The idea is that as long as mortgages are closely tied to short-term interest rates, unexpected increases in the policy rate will quickly lead to a shift in cash flows and mortgage payments, which is particularly important for existing borrowers. In this sense, changes in the policy rate also affect the initial cost of new home loans, which in turn affects the demand for housing. The ARM share in the USA exhibits a decline since the early 1980s, where the ARM share data starts. This encourages us to employ the dynamic stochastic general equilibrium (DSGE) model of Alpanda and Zubairy (2017) and examine exogenous variation in the ARM share. The model features a housing market and an occasionally binding credit constraint along the lines of Iacoviello (2005). The impact of the ARM share is nonlinear: a drop in the ARM share causes a less than proportional decline in the interest rate elasticity of mortgage debt. We calibrate the model to the ARM share in 1985, leaving all other parameters at their sample average, and to the ARM share at the end of the data sample. The resulting responses of mortgage debt to a simulated 25bp tightening shock quantitatively matches the estimated decline in the mortgage response sensitivity to policy shocks.

Our results have important policy implications. First, a weaker transmission of policy impulses to the mortgage market, but not to the real economy, implies that monetary policy is not the right instrument to reduce household sector debt. A policy tightening at the end of the sample period to curb the build-up of mortgage debt is expensive in terms of forgone employment. Hence, we would advise against using monetary policy to counteract financial risks related to household

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

borrowing. This finding is in line with Alpanda and Zubairy (2017) who find that, for example, maximum loan-to-value (LTV) ratios are better macroprudential tools to counteract high household indebtedness.

Second, the results can be interpreted as a case against the "too low for too low" argument. That is, keeping interest rate low for an extended period of time does not contribute to inflating a housing credit boom. Our counterfactuals show that the systematic component of monetary policy contributes little to the time-varying sensitivity of mortgage debt to monetary policy.

Our results are related to several branches of the literature. In the following, we highlight only those papers which we consider most relevant for us. The first field studies structural features of the US mortgage market and relates them to the strength of the transmission process. Calza et al. (2013) present VAR results consistent with the notion that the monetary transmission to housing investment is stronger for a high share of ARM. We extend this line of research by looking at the US economy over time, not at the cross section of countries which differ in the average ARM share. Zeev (2016) presents a partial equilibrium model of the housing market, in which a high share of ARM mortgage contracts amplifies the effect of an interest rate shock. He also presents empirical evidence consistent with this finding. The ARM share is used to interact the economy's response to a credit supply shock. Garriga et al. (2017) build a general equilibrium model with incomplete asset markets in which monetary policy affects housing investment by changing the cost of new mortgages. A high ARM share again intensifies the response of the economy to the monetary policy impulse.[4] Secondly, the result of this paper can be interpreted as an equivalent to the "mortgage rate conundrum" of Justiniano et al. (2017). They establish the finding that the connection between mortgage interest rates and Treasury rates broke down in 2003. Hence, the policy tightening of the Fed in 2004 did not lead to higher mortgage rates. Paul (2020) also uses a TVP-VAR model to study the policy transmission to various asset prices. He finds that the transmission to house prices was particularly weak before the 2008–2009 financial crisis.

The remainder of this paper is structured as follows: section two sketches the evolution of mortgage debt. Section three introduces the TVP-VAR model and details about the Bayesian estimation. Section four discusses the main results. Section five focuses on the role of adjustable rate mortgages for our results and presents results from a DSGE model. Sections six and seven present counterfactual analyses and robustness checks, respectively. Section eight concludes. An online appendix contains additional information.

2. Mortgage debt

This paper focuses on mortgage debt held by US private households, which is the main component of overall household debt. The left panel in Figure 1 depicts the development of postwar real household debt. There has been a steady increase in household debt and mortgage debt, respectively, until the eve of the Great Recession.

For the purpose of this paper, we look at the cyclical component of real mortgage debt, which we derive from Baxter and King (1999) filtering the original series. The filter has a band length of eight and lets frequencies pass that are between 4 and 64 quarters. With this calibration, we account for the average length of financial and debt cycles in the USA and take into account that these cycles are about twice as long as the business cycle (Alpanda and Zubairy, 2019). The resulting cyclical series, the mortgage gap, is depicted in the right panel of Figure 1.

We find that the mortgage gap fluctuates between 4% and −5% relative to its trend and peaks before each recession. During a recession, the gap turns negative.

3. The empirical model

We rely on a VAR model with drifting coefficients as well as a drifting variance–covariance matrix. To introduce a few key elements and fix notation, we start with a time-invariant VAR model.

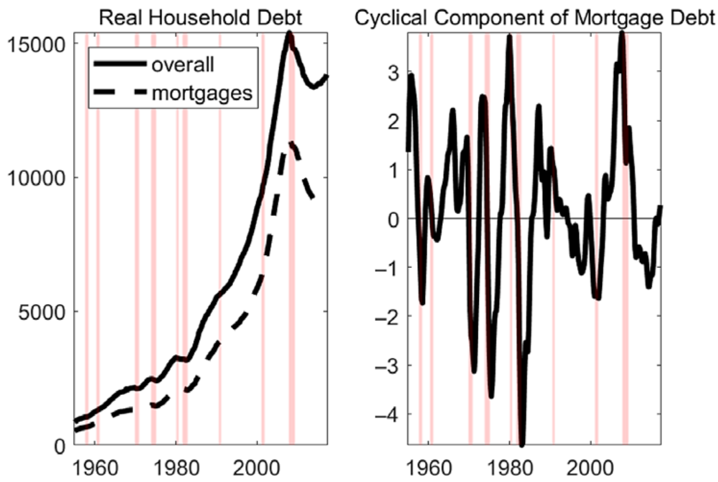

Figure 1. Household Debt in the USA.
Notes: The solid line in the left panel is overall real household debt. The dashed line in the left panel shows real mortgage debt. The right panel shows the cyclical component of Baxter–King-filtered mortgage debt (in %). The shaded areas reflect NBER-dated recessions.

3.1 Structural VAR model

A standard time-invariant structural VAR is defined as:

$$\mathbf{A}\mathbf{y}_t = \mathbf{d} + \mathbf{F}_1\mathbf{y}_{t-1} + ... + \mathbf{F}_s\mathbf{y}_{t-s} + \boldsymbol{\varepsilon}_t, \qquad t = s+1, ..., T, \tag{1}$$

where $\mathbf{y}_t$ is a $k \times 1$ vector of observed endogenous variables, $\mathbf{d}$ is a $k \times 1$ vector which includes deterministic components, and $\mathbf{A}, \mathbf{F}_1, ..., \mathbf{F}_s$ are $k \times k$ matrices of VAR coefficients. Finally, $\boldsymbol{\varepsilon}_t$ is a $k \times 1$ vector of structural shocks, where $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(0, \boldsymbol{\Sigma}\boldsymbol{\Sigma}')$.

The vector $\mathbf{y}_t$ contains four variables. The first variable is the unemployment rate, which is our measure of real economic activity. The second variable is inflation, measured as the year-on-year growth rate of the GDP deflator. The third variable is the cyclical component of real mortgage debt derived before. Our fourth variable is supposed to reflect the Fed's policy instrument. We therefore use the effective federal funds rate, augmented with the Wu and Xia (2016) shadow rate during periods in which the federal funds rate was characterized by the zero lower bound (ZLB).[5] Note that our choice of variables other than the cyclical component of mortgage debt is almost the same as in Primiceri (2005).[6] We refer to the first three variables as the non-policy block of the VAR system, for reasons to become clear below.

The simultaneous relations of structural shocks are specified by recursive identification, assuming that $\mathbf{A}$ is lower triangular as:[7]

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ a_{21} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ a_{k1} & \cdots & a_{kk-1} & 1 \end{bmatrix}. \tag{2}$$

Premultiplying both sides by $\mathbf{A}^{-1}$, the model in (1) can hence be rewritten as:

$$\mathbf{y}_t = \mathbf{c} + \mathbf{B}_1\mathbf{y}_{t-1} + \cdots + \mathbf{B}_s\mathbf{y}_{t-s} + \mathbf{A}^{-1}\boldsymbol{\Sigma}\boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(0, \mathbf{I}_k), \tag{3}$$

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

where $\mathbf{c} = \mathbf{A}^{-1}\mathbf{d}$ and $\mathbf{B}_j = \mathbf{A}^{-1}\mathbf{F}_j$ for $j = 1, ..., s$. The standard deviations of the structural shocks are captured in the $k \times k$ matrix $\mathbf{\Sigma}$ which is specified as:

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sigma_k \end{bmatrix}. \tag{4}$$

Stacking all elements of $\mathbf{c}$ and $\mathbf{B}_j$, we get the $k + (k^2 s) \times 1$ vector $\mathbf{B}$. Defining $\mathbf{X}_t = \mathbf{I}_k \otimes [1, \mathbf{y}'_{t-1}, ..., \mathbf{y}'_{t-s}]$, the model can be rewritten in reduced-form as:

$$\mathbf{y}_t = \mathbf{X}_t \mathbf{B} + \mathbf{A}^{-1} \mathbf{\Sigma} \boldsymbol{\varepsilon}_t. \tag{5}$$

We order the unemployment rate in $\mathbf{y}_t$ first, the inflation rate second, mortgage debt third, and the effective federal funds rate last. This ordering of variables implies that all variables of the non-policy block belong to the information set of the central bank, such that we allow the central bank to react contemporaneously to these variables. At the same time, our timing restriction implies that the monetary policy shock, which we are primarily interested in, does not contemporaneously affect the non-policy block. That is, we assume that unemployment, inflation, and mortgage debt react to monetary policy shocks with at least one period delay. While this assumption is standard in the literature (see Christiano et al., 1999) for unemployment and inflation, it is also reasonable to assume that mortgage debt does not contemporaneously respond to unexpected changes in the federal funds rate.

Note that including mortgages into a VAR model also implies that, in principle, the policy reaction function incorporated in the VAR model allows for a feedback from mortgages to monetary policy. Hence, the Fed could respond to changes in household indebtedness. Since we want to keep the model as close as possible to the standard VAR model, even when including mortgages, we restrict the response of monetary policy to mortgages to zero across all lags.[8]

The model is estimated on quarterly data from 1955Q3 to 2017Q2. We use $s = 2$ two lags. This choice comprises a reasonable trade-off between the risk of over-parameterization of our model on the one hand, and autocorrelated residuals on the other.[9]

3.2 The TVP-VAR model

Our TVP-VAR is specified as:

$$\mathbf{y}_t = \mathbf{X}_t \mathbf{B}_t + \mathbf{A}_t^{-1} \mathbf{\Sigma}_t \boldsymbol{\varepsilon}_t, \tag{6}$$

where all parameters, that is, the VAR coefficients captured in $\mathbf{B}_t$, the simultaneous relationships among the endogenous variables captured in $\mathbf{A}_t$ as well as the stochastic volatility of our structural shocks captured in $\mathbf{\Sigma}_t$, are time-varying. Let $\mathbf{a}_t = [a_{21,t}, \cdots, a_{kk-1,t}]'$ be the vector of nonzero and non-one elements in $\mathbf{A}_t$ (i.e., the lower-triangular elements in $\mathbf{A}_t$) and let $\mathbf{h}_t = [h_{1,t}, \cdots, h_{k,t}]'$ with $h_{i,t} = \log \sigma^2_{i,t}$ for $i = 1, \cdots, k$. Following Primiceri (2005) and others, we assume the following dynamics of the model's parameters

$$\mathbf{B}_{t+1} = \mathbf{B}_t + \eta_{B,t}, \tag{7}$$

$$\mathbf{a}_{t+1} = \mathbf{a}_t + \eta_{a,t}, \tag{8}$$

$$\mathbf{h}_{t+1} = \mathbf{h}_t + \eta_{h,t}. \tag{9}$$

As regards the initial states of the time-varying parameters, we follow Nakajima (2011) and assume that $\mathbf{B}_{s+1} \sim \mathcal{N}(\mu_{B_0}, \mathbf{\Sigma}_{B_0})$, $\mathbf{a}_{s+1} \sim \mathcal{N}(\mu_{a_0}, \mathbf{\Sigma}_{a_0})$ and $\mathbf{h}_{s+1} \sim \mathcal{N}(\mu_{h_0}, \mathbf{\Sigma}_{h_0})$. Finally, we follow

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

Primiceri (2005) and assume that all innovations in the model are jointly normally distributed as:

$$\mathbf{V} = \mathrm{Var}\left(\begin{bmatrix} \varepsilon_t \\ \eta_{B,t} \\ \eta_{a,t} \\ \eta_{h,t} \end{bmatrix}\right) = \begin{bmatrix} \mathbf{I}_k & 0 & 0 & 0 \\ 0 & \boldsymbol{\Sigma}_B & 0 & 0 \\ 0 & 0 & \boldsymbol{\Sigma}_a & 0 \\ 0 & 0 & 0 & \boldsymbol{\Sigma}_h \end{bmatrix}. \tag{10}$$

Note that the parameters follow a random walk process. Although the random walk process is nonstationary, its assumption can capture both gradual and sudden structural changes[10]. We follow the literature and choose the random walk assumption because we do not want to impose further restrictions on the transition equation a priori when estimating our state-space model. Moreover, the advantage of the random walk assumption is that (1) it allows for permanent shifts and (2) drastically reduces the number of parameters in the estimation procedure.[11]

3.3 Estimation algorithm and priors

In order to estimate the TVP-VAR, we rely on MCMC methods. Our estimation procedure mainly follows Nakajima (2011) and can be summarized as follows: given the data $\mathbf{y} = \{\mathbf{y}_t\}_{t=1}^{T}$, $\boldsymbol{\omega} = (\boldsymbol{\Sigma}_B, \boldsymbol{\Sigma}_a, \boldsymbol{\Sigma}_h)$ and our prior density $\pi(\boldsymbol{\omega})$, we use the following MCMC algorithm to draw samples from the posterior $\pi(\mathbf{B}, \mathbf{a}, \mathbf{h}, \boldsymbol{\omega}|\mathbf{y})$[12]:

(1) Initialize $\mathbf{B}$, $\mathbf{a}$, $\mathbf{h}$ and $\boldsymbol{\omega}$.
(2) Sample $\mathbf{B} \mid \mathbf{a}, \mathbf{h}, \boldsymbol{\Sigma}_B, \mathbf{y}$.
(3) Sample $\boldsymbol{\Sigma}_B \mid \mathbf{B}$.
(4) Sample $\mathbf{a} \mid \mathbf{B}, \mathbf{h}, \boldsymbol{\Sigma}_a, \mathbf{y}$.
(5) Sample $\boldsymbol{\Sigma}_a \mid \mathbf{a}$.
(6) Sample $\mathbf{h} \mid \mathbf{B}, \mathbf{a}, \boldsymbol{\Sigma}_h, \mathbf{y}$.
(7) Sample $\boldsymbol{\Sigma}_h \mid \mathbf{h}$.
(8) Go back to 2.[13]

Priors need to be specified for the starting values of our MCMC algorithm (i.e. for the initial state of the time-varying parameters $(\mu_{B_0}, \boldsymbol{\Sigma}_{B_0})$, $(\mu_{a_0}, \boldsymbol{\Sigma}_{a_0})$ and $(\mu_{h_0}, \boldsymbol{\Sigma}_{h_0})$) and for the i^{th} diagonals of the covariance matrices (i.e., the i^{th} diagonals of $\boldsymbol{\Sigma}_B$, $\boldsymbol{\Sigma}_a$ and $\boldsymbol{\Sigma}_h$). As regards the initial states for the VAR parameters, there are mainly two common practices for specifying the initial states. The first approach follows Primiceri (2005), who uses the point estimates of a time-invariant VAR model estimated from, say, the first 10 years, in order to calibrate the prior. The potential drawback of this approach is that we lose these observations for the estimation of our TVP-VAR, as in a standard Bayesian setting the prior should not contain any information based on the sample. Second, from the standpoint that we do not have any information about the initial state a priori, setting diffuse priors for the initial states is another option (see Nakajima, 2011, for a discussion). In particular, the initial state of our parameters has flat priors, set as:

$$\mathbf{B}_{s+1} \sim \mathcal{N}(0, 10 \cdot I), \qquad \mathbf{a}_{s+1} \sim \mathcal{N}(0, 10 \cdot I), \qquad \mathbf{h}_{s+1} \sim \mathcal{N}(0, 10 \cdot I). \tag{11}$$

Of course, the prior choice for the hyperparameters can affect posterior inference, although $\boldsymbol{\Sigma}_B$, $\boldsymbol{\Sigma}_a$, and $\boldsymbol{\Sigma}_h$ do not parameterize time variation in the first line, but only prior beliefs about the time variation. However, as emphasized by Primiceri (2005), Nakajima (2011), and Koop and Korobilis (2010), the priors should be carefully chosen because our model has many parameters to estimate. This is even more important from the standpoint that all of our model's parameters are modeled as nonstationary random walks. Thus, tight priors for the variance–covariance matrices of the disturbances in the random walk process should avoid ill-determined (implausible)

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

behaviors of the parameters. Moreover, it should be noted that in general the time-varying VAR coefficients require a tighter prior than the time-varying variance–covariance matrix. To address this problem, we rely on rather tight priors.

In particular, we set the following priors for the i^{th} diagonals of $\mathbf{\Sigma}_B$, $\mathbf{\Sigma}_a$, and $\mathbf{\Sigma}_h$:[14]

$$(\mathbf{\Sigma}_B)_i^{-2} \sim \Gamma(25, 6 \cdot 10^{-5}), \qquad (\mathbf{\Sigma}_a)_i^{-2} \sim \Gamma(4, 0.05), \qquad (\mathbf{\Sigma}_h)_i^{-2} \sim \Gamma(4, 0.05).$$

To compute the posterior estimates, we draw $n = 20,000$ draws and discard the first 10, 000 draws, as samples that have been generated in early iteration steps are likely to be not representative for the true posterior distribution. The appendix reports several convergence checks that are common in the literature, including inefficiency factors as well as the convergence checks proposed by Geweke (1992) and Raftery and Lewis (1992). It stands out that all convergence checks justify our prior choice and that our Markov chain mixes quickly.

We further investigate whether the time variation found in our coefficients is a feature of the data or driven by our priors. For that purpose, we apply the test proposed by Cogley and Sargent (2005) and compare the trace of the prior with the trace of the posterior distribution of $\mathbf{\Sigma}_\beta$. We would point to time variation in the parameters stemming from the data rather than the priors if the trace of the prior is below the 16^{th} percentile of the trace of posterior distribution. As can be seen in the appendix, the trace of the prior is indeed well below the 16^{th} percentile of the trace of the posterior of $\mathbf{\Sigma}_\beta$.

4. Results

The advantage of our TVP-VAR model is that we can show time-varying effects of monetary policy shocks. This time variation is not only driven by the estimated VAR coefficients in $\mathbf{B}_t$, but also by the shock impact matrix, $\mathbf{A}_t^{-1}$, as well as the stochastic volatility of the covariance matrix, $\mathbf{\Sigma}_t$.[15]

4.1 Responses to a monetary policy shock

In this section, we discuss the time-varying effects of monetary policy shocks on the endogenous variables. To get an impression of the changing nature of monetary policy transmission that accounts for estimation uncertainty, Figure 2 plots the impulse responses following monetary policy shocks originating in 1985Q1 and 2011Q3, respectively. These dates are chosen among episodes of a high and low ARM share, which will become important in the next section. The left column shows the impulse responses following a 25bp monetary policy shock in 1985Q1, when the ARM share was about 60%, whereas the middle column shows the corresponding impulse responses in 2011Q3, when the ARM share reached a minimum of about 10%.

An unexpected policy tightening in 1985 raises unemployment by 0.05 percentage points and leads to a significant fall in mortgage debt by about 0.03%. The response of the inflation rate exhibits a price puzzle, that is, a positive response in the first few quarters, before it starts to fall. A shock in 2011 triggers a similarly sized increase in unemployment and a drop in inflation. However, mortgage debt is no longer responsive to monetary policy.

The right panel shows the difference between the impulse responses in 1985Q1 and 2011Q3, respectively. We would refer to a significantly different transmission of monetary policy shocks when the error bands do not enclose the zero line. Contrary to this, the transmission would not be significantly different across the two points in time when the the zero line lies inside the error bands.

For mortgages, the impulse response is significantly stronger in the early 1980s than toward the end of our sample. The responses of unemployment and inflation, however, are not significantly different across the two points in time. Hence, the sensitivity of mortgage debt to monetary policy

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

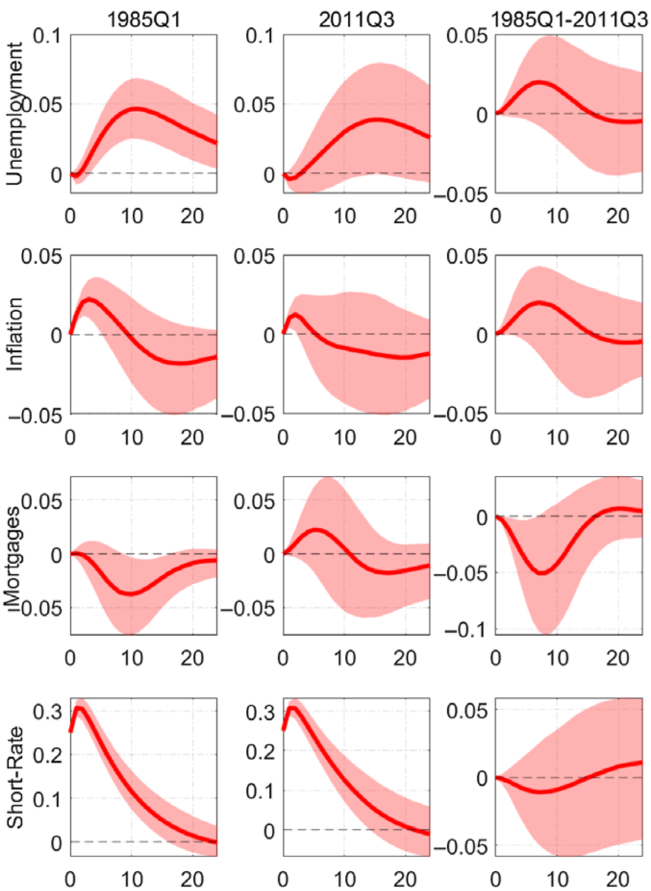

Figure 2. Response to a Monetary Policy Shock for Selected Dates.

Notes: The left panel shows the mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock in 1985Q1, whereas the middle panel shows the mean responses to a 25bp monetary policy shock in 2011Q3. The right panel shows the difference between the responses in 1985Q1 and 2011Q3.

shocks is lower at the end of our sample period, while the transmission of policy to inflation and unemployment remains unaffected. This is the most important finding of this paper.

Notice that this finding does not only hold for the two selected dates, that is, for 1985Q1 and 2011Q3. Rather, we find that the transmission of monetary policy shocks to mortgage debt gradually decreases. We best see that when we study the cross sections of the time-varying impulse response functions with their respective 16^{th} and 84^{th} percentiles of the posterior distribution.

Figures 3, 4, and 5 show the time-varying reaction of the endogenous variables 4, 8, and 16 quarters after the shock. The full profiles of impulse responses, that is, the response as a function of the timing of the occurrence of the shock and the time after the shock, are shown as three-dimensional plots in the appendix.

After four quarters, see Figure 3, the policy tightening has led to higher unemployment. The response of unemployment is significant between the 1970s and 1990s. For inflation, we observe a price puzzle in the 1970s. However, the initial significance of our observed price puzzle begins to vanish over time. Mortgage debt falls significantly following a policy shock originating in the 1960s and 1970s but remains unchanged for shocks occurring more recently.

Eight quarters after the shock, see Figure 4, unemployment is still higher as a result of the tightening, though the impact on unemployment becomes slightly smaller over time. The response of inflation 2 years after the shock remains insignificant for most of the sample period. Importantly, we see a very pronounced tendency in the response of mortgage debt: the sensitivity falls steadily over time and disappears in the 2000s.

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

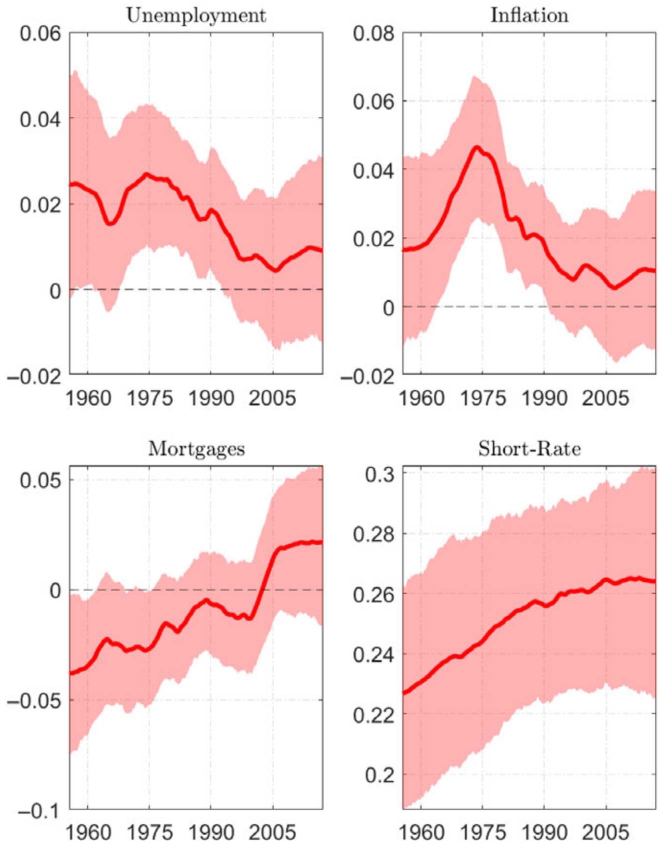

Figure 3. Response to a Monetary Policy Shock After Four Quarters.
Notes: Mean responses (red-solid) after four quarters with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

The impact of the policy shock after 4 years is shown in Figure 5. Again, the response of mortgage debt is negative between the 1970s and the 1990s. After that, mortgages remain insignificantly different from their trend. From a first glance at Figure 5, one might mistakenly infer that the mortgage response will regain strength over time. However, this is due to the fact that the response to a monetary policy shock becomes somewhat more sluggish over time with the peak response occurring later. Thus, especially for the last 20 years of our sample, Figure 5 shows the response of mortgages at the time when the response to monetary policy is strongest, before mortgages return to the trend.

The persistence of the interest rate response to the monetary policy shock does not fluctuate much over time.

Finally, we extract the maximum impact of the monetary policy shock on mortgage debt as well as the time at which the maximum response occurs and show the evolution of both in Figure 6.[16] This allows us to better describe the shift in the sensitivity of mortgage debt to monetary policy over the sample period. We can see the general tendency of diminishing effects of monetary policy shocks with responses becoming insignificant at the end of our sample. We also find that the maximum impact of the policy shock occurs later over time.

The decline in the sensitivity of mortgage debt has strong implications for monetary policy, which we will now study in detail. Keeping in mind that the reduced-form model is not suitable for normative policy analysis, our results suggest that using monetary policy to curb household debt is difficult. This is because the effectiveness of monetary policy shocks strongly varies over

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

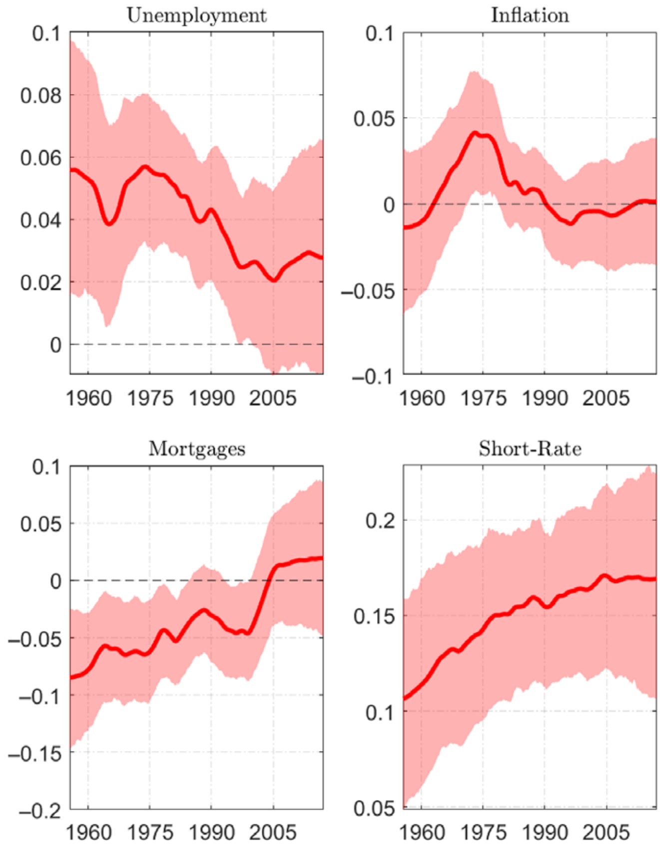

Figure 4. Response to a Monetary Policy Shock After Eight Quarters.
Notes: Mean responses (red-solid) after eight quarters with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

time. This also implies that the costs of a debt reduction in terms of unemployment vary over time. We will pick up this thought in the following section.

4.2 The sacrifice ratio of debt reduction

Let us consider a central bank which aims at using an increase in the policy rate in order to foster a debt reduction of households. How much additional unemployment does the central bank need to accept? We invoke the concept of the sacrifice ratio which in its initial form answers the question of how costly a disinflation is in terms of unemployment. We construct a measure that shows how expensive a debt reduction is in terms of unemployment. This measure, $\varphi_{t,t+h}$, is the ratio of the cumulative response of mortgage debt relative to the cumulative response of unemployment, h periods after the shock occurs. That is, the nominator and the denominator are the cumulative mean impulse responses:

$$\varphi_{t,t+h} = \frac{\sum_{i=0}^{h} \mathrm{IRF}_{t+i}^{\mathrm{debt}}}{\sum_{i=0}^{h} \mathrm{IRF}_{t+i}^{\mathrm{unemp}}}, \tag{12}$$

where the horizon h can be interpreted as the relevant horizon for policy-makers. Consider a policy shock that leads to a reduction in debt and an increase in unemployment. Hence, the ratio is negative. If for the same loss in employment debt falls more, this reduction becomes less costly. In this case, the index falls. By the same token, we would observe a drop in the sacrifice ratio

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

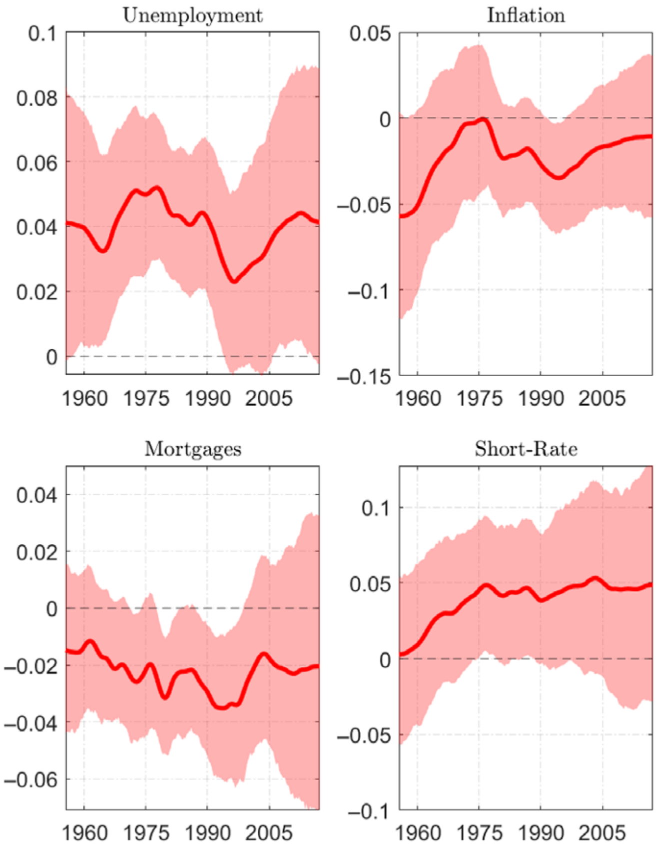

Figure 5. Response to a Monetary Policy Shock after 16 Quarters.
Notes: Mean responses (red-solid) after 16 quarters with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

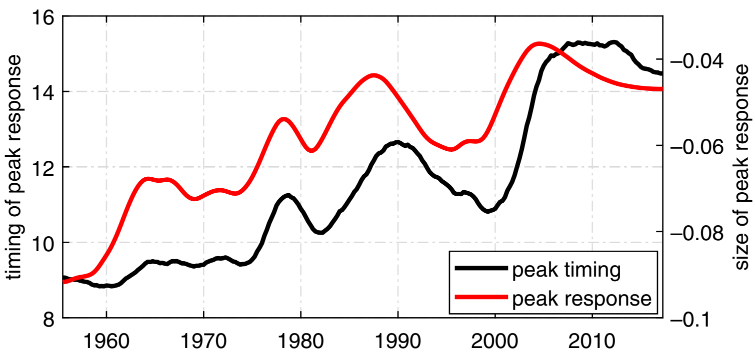

Figure 6. Peak Response of Mortgage Debt.
Notes: Peak response and timing of the peak response following a 25bp monetary policy shock.

if unemployment increases less, given that the reduction of debt remains unchanged. That is, policy-makers were able to achieve a high reduction in mortgage debt per unit increase in unemployment. An increase in $\varphi_{t,t+h}$ would thus be consistent with a more costly reduction, that is, for a one unit increase in unemployment we only get a relatively small reduction in debt. We consider two alternative horizons: $h=8$ and $h=12$ quarters. Since we base this thought experiment on a

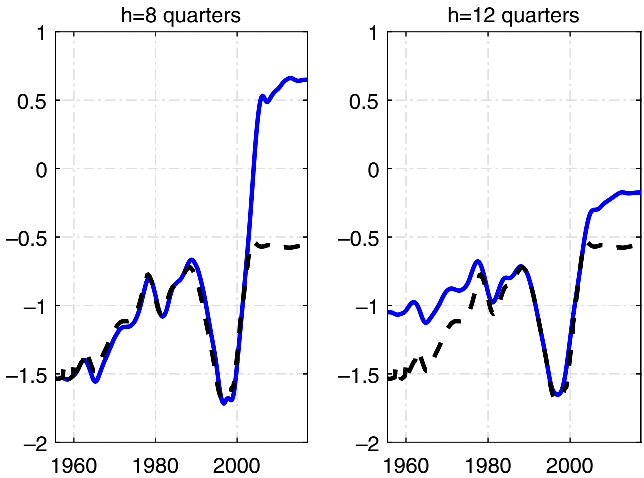

Figure 7. The Costs of Debt Reduction.
Notes: The blue-solid lines reflect the ratio $\varphi_{t,t+h}$ for horizon $h = 8$ (left panel) and $h = 12$ (right panel) over time. In both panels, the black-dashed line reflect $\varphi_{t,t+h}$ where h corresponds to the timing of the peak response.

reduced-form model, we should stress that we do not intend to derive normative implications. Rather, we want to find out at what time a debt reduction as a result of a monetary tightening was relatively expensive or cheap, respectively.

The blue-solid lines in Figure 7 display the ratio for the two horizons h, whereas the black-dashed lines plot $\varphi_{t,t+h}$ with h corresponding to the timing of the peak response. We can see an upward trend in both series. Hence, a debt reduction becomes more expensive over time. The trend is broken only during the 1990s.

Technically, the break in the sacrifice ratio results from the changes in the impulse responses of unemployment and mortgages in the first half of the 1990s. The response of unemployment eight quarters after the shock remains horizontal at that time, while the response of mortgages temporarily become somewhat more pronounced. One possible explanation for the behavior of the unemployment response is the sudden weakening of Okun's Law (Owyang and Sekhposyan, 2012). Hence, when the Fed tightened, both mortgage debt and real economic activity fell. Due to the shift in Okun's Law, however, the loss in terms of employment was smaller than before and the debt reduction appeared more favorably when compared to the loss in employment. The sudden shift in the responsiveness of mortgages can best be explained in terms of changes in the ARM share. As shown in below, the ARM share jumped from 20% to more than 50% in the early- to mid-1990s, before it continued its downward trend. This jump is directly reflected in a stronger response of mortgages to monetary policy shocks.

Overall, our results mirror-image the time-varying sensitivity of unemployment and debt, that is, that monetary policy-induced debt reductions are (relatively) cheap if either the sensitivity of unemployment to monetary policy shocks is relatively low while the sensitivity of mortgage debt is relatively high or vice versa.

5. The role of adjustable rate mortgages

Many structural forces could potentially be responsible for the decline in the sensitivity of mortgages to Fed policy. Calza et al. (2013) highlight key parameters that have an effect on how sensitive mortgages are to monetary policy. They find that the transmission of monetary policy shocks to consumption is stronger in countries where mortgages with adjustable rates dominate. In the US economy, this share declined over time. Figure 8 plots the ARM share since the early 1980s. While in 1985 the share was about 60%, it fell to only 10% by the end of 2011, implying that home-buyers mostly prefer fixed-rate mortgages over adjustable rate mortgages.[17] As emphasized

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

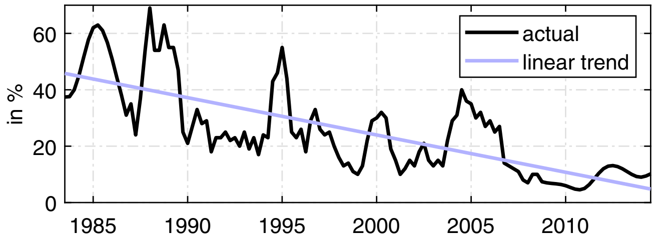

Figure 8. Share of Adjustable Rate Mortgage Contracts.
Notes: The data are taken from Federal Housing Agency, Monthly Interest Rate Survey. We use chained data based on data availability. Thus, we use interpolated data extracted from annual basis from 1982Q2 until 1984Q4 and end-of-quarter monthly data from 2008Q4 until 2014Q3.

by Alpanda and Zubairy (2017), under fixed-rate contracts the interest rate channel of monetary policy can be impaired, as agents are unable to lower the interest burden on their existing debt by refinancing, especially when their existing debt levels are high and they do not possess adequate equity. We therefore expect that the transmission of monetary policy shocks to the economy is linked to the ARM share. In particular, we expect that the transmission is more powerful when the ARM share is high.

With this in mind, let us briefly return to Figure 2. We show the impulse responses for shocks originating at a time when the ARM share is high (1985Q1) and low (2011Q3), respectively. The two impulse response are indeed significantly different.

5.1 Illustrative evidence

We want to obtain some illustrative evidence on the determinants of time variation in the policy impact on mortgages. For that purpose, we regress the dynamic impulse response functions on the ARM share. We also control for other characteristics of the prevailing contracts in this market.[18] This is done for both the time-varying impulse responses and the cumulative effects of mortgage debt for different periods h after the shock. Formally, we regress

$$\mathrm{IRF}_{t,t+h} = c + \boldsymbol{\gamma}\mathbf{x}_t + \delta t + \boldsymbol{\varepsilon}_t,$$

as well as for impulse responses accumulated up to horizon h,

$$\sum_{i=0}^{h} \mathrm{IRF}_{t,t+i} = c + \boldsymbol{\gamma}\mathbf{x}_t + \delta t + \boldsymbol{\varepsilon}_t,$$

where t is the period the shock occurs and i denotes the number of periods after the shock.

The vector $\boldsymbol{\gamma}$ collects the coefficients for the control variables in $\mathbf{x}_t$, c is a constant, and $\boldsymbol{\varepsilon}_t$ is an error term. The matrix $\mathbf{x}_t$ contains the ARM share, the annual growth rate of GDP, the S&P/Case-Shiller US National Home Price Index (in logs), the LTV ratio, and the average effective interest rate payed for mortgage contracts (effective rate), which is also provided by the Federal Housing Agency's Interest Rate Survey. To account for possible nonlinearities between the impulse responses and the ARM share, $\mathbf{x}_t$ also includes the squared ARM share. This is primarily motivated by the implications of our DSGE model, which will be described in the following section. Lastly, we add a linear trend, which is meant to account for the fact that both the dependent variable and some of the explanatory variables are very persistent. Heteroskedasticity- and autocorrelation-consistent (HAC) standard errors are used with a lag length of eight, as recommended by Wooldridge (2016).

Table 1 shows the outcome of the regressions for $h = 8, 12, 16$ and h being equal to the timing of the peak response. A few things stand out. First, for all horizons under consideration, the joint explanatory power of our control variables is high and in all but one case well above 85%. Most control variables also have a significant impact and can explain some of the variation in the time-varying effects of mortgages. Second, the coefficient for the ARM share shows the expected negative sign in all cases and is significant in almost all cases except for $h = 8$ in the upper panel as well as for the peak response in the bottom panel. This implies that mortgage debt falls more

Table 1. Explaining the Time-Varying Policy Impact on Mortgages.

	Impulse Response			
	Peak	$h = 16$	$h = 12$	$h = 8$
Constant	0.025 (0.039)	0.010 (0.021)	0.052 (0.071)	0.054 (0.123)
ARM share	−0.349** (0.158)	−0.207* (0.119)	−0.452** (0.231)	−0.580 (0.370)
ARM share squared	0.706*** (0.175)	0.366** (0.147)	0.948*** (0.261)	0.125*** (0.039)
Real GDP	0.530* (0.317)	0.995** (0.396)	0.594 (0.544)	−0.192 (0.645)
House prices	0.030*** (0.006)	0.012*** (0.001)	0.022** (0.009)	0.026 (0.021)
LTV	−2.706*** (0.617)	−1.466*** (0.255)	−2.99*** (0.661)	−3.157*** (0.751)
Effective rate	1.065* (0.542)	1.818*** (0.298)	2.873*** (0.746)	2.741* (1.344)
Trend	yes	yes***	yes**	yes
R^2	0.877	0.683	0.878	0.875
adj. R^2	0.817	0.664	0.870	0.867
	Cumulative Impulse Response			
	Peak	$h = 16$	$h = 12$	$h = 8$
Constant	0.425 (0.975)	0.640 (1.204)	0.550 (1.133)	0.285 (0.683)
ARM share	−3.410 (2.568)	−6.263* (3.558)	−5.194* (3.261)	−3.015* (1.785)
ARM share squared	0.696** (0.297)	0.129** (0.376)	1.080*** (0.342)	0.622*** (0.204)
Real GDP	2.463 (3.462)	2.273 (6.355)	−0.674 (5.234)	−1.993 (3.145)
House prices	0.127 (0.151)	0.285 (0.185)	0.216 (0.194)	0.113 (0.123)
LTV	−19.955*** (0.610)	−35.496*** (8.362)	−27.583*** (6.676)	−13.967*** (3.553)
Effective rate	26.151** (8.424)	26.151*** (8.424)	26.395** (11.427)	12.966* (6.995)
Trend	yes	yes	yes	yes
R^2	0.872	0.888	0.878	0.872
adj. R^2	0.865	0.881	0.870	0.865

Notes: The dependent variable is the response of mortgage debt following a 25bp monetary policy shock when the response peaks (first column), after 16 periods (second column), after 12 periods (third column), and after 8 periods (fourth column). The coefficients and standard errors of the ARM share, real GDP, the loan-to-value ratio, the short-rate, and trend are multiplied by 1000. The coefficients and standard errors of the squared ARM share are multiplied by 10,000. HAC standard errors are used with a lag length of eight and a Newey–West fixed bandwidth of 5. A significance level of 1%, 5%, and 10% is denoted by ***, **, and *.

strongly after a policy tightening when the ARM share increases, which is in line with the results found in Alpanda and Zubairy (2017), Calza et al. (2013), and Garriga et al. (2017).

Second, the coefficient on the squared share of adjustable rate mortgages is positive and significant in all cases. Note that a positive coefficient implies that the effect of an increase in the ARM share is decreasing, the higher the ARM share is. This means that the response of mortgages after an increase in the ARM share is stronger in absolute terms when the ARM share is relatively low. If, on the other hand, the ARM share is relatively high, a further increase in the ARM share will still lead to a stronger response, this effect is however less pronounced than before.

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

Third, the coefficient on the LTV ratio is negative and significant in all cases. This result is in line with Alpanda and Zubairy (2017) and implies that the effects of monetary policy shocks are amplified when the LTV ratio increases.

Overall, our results support our assumption that the ARM share affects the transmission of monetary policy shocks.

5.2 A DSGE model with mortgage contracts

While the previous regressions are illustrative, eventually a structural model is needed to shed light on the impact of a shift in the ARM share on the strength of policy transmission. Therefore, we resort to the DSGE model by Alpanda and Zubairy (2017).

The model builds on a closed-economy DSGE with housing and household debt as well as an occasionally binding credit constraint. The model features two types of households: patient households (savers) and impatient households (borrowers). In the model, excessive household debt arises as a result of exuberance shocks on expectations on house prices, which drives a wedge between actual and fundamental values.

Importantly, the model allows the average duration of the fixed interest rate for loans to be shorter than the full amortization duration of the underlying loan itself. That is, the interest rate on new mortgage loans is decomposed into a fraction carrying a fixed mortgage interest rate and a fraction of existing loans that is refinanced each period.

We use the model to simulate impulse responses to monetary policy shocks for different calibrations of the ARM share. We simulate impulse responses for mortgage debt to a 25bp monetary policy shock, which is consistent with the definition of the policy shock in the TVP-VAR. In the first case, the "low share" case, we use the same overall calibration as Alpanda and Zubairy (2017), including the interest rate adjustability of mortgages based on a 10% ARM share, which we can observe at the end of our sample. In the second case, the "high share" case, we recalibrate the interest rate adjustability parameter based on an ARM share of 60% as observed in 1985, keeping anything else similar to the benchmark case.[19]

In Figure 9, we compare the impulse responses to 25bp monetary policy shock for both the DSGE model and the TVP-VAR. First, similar to our time-varying VAR, there is a weaker reaction of mortgage debt to monetary policy shocks when the ARM share is low. Second, the amplitude of both the DSGE and the TVP-VAR are nearly identical in the high share case with a 0.05 percentage point drop in mortgages. For the low share case, the TVP-VAR shows a weaker response. However, this can most likely be attributed to the fact that our recalibration was solely based on the different fraction of ARM shares keeping all other parameters unchanged. In the model, the interest rate duration is based on factors other than the ARM share, including home equity loans and repayments. Additionally, we can see that the relationship between the ARM share and the size of the mortgage reaction to a restrictive 25bp monetary policy shock shows a nonlinear pattern.

To conclude this section, we calibrate the DSGE model not just for a high and a low share, but for the full range of observable ARM shares. Figure 10 plots the resulting peak responses to the monetary policy shock against the ARM share. The sensitivity to the policy shock increases nonlinearly in the ARM share, which is why we allow for a nonlinear effect in our regressions reported in Table 1. Nevertheless, although we also find a nonlinear relationship here, the sign in this case does not match the sign from the illustrative regression results.

Summing up, the DSGE model provides further evidence for the role of ARM shares in the transmission of monetary policy shocks. The empirically observed drop in the ARM share since the early 1980s leads to impulse responses which are quantitatively very similar to the responses derived from the TVP-VAR.

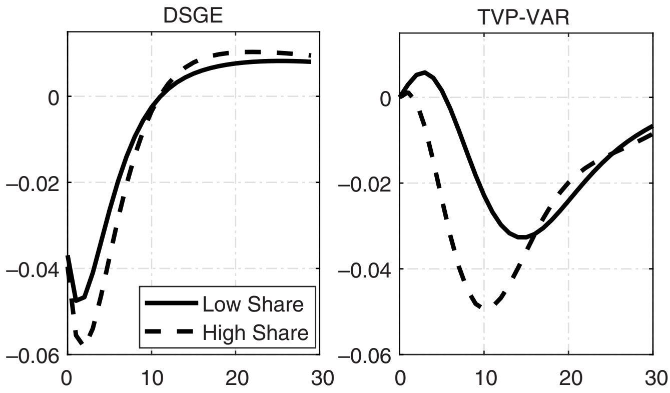

Figure 9. Responses of Mortgage Debt to Monetary Policy in the DSGE Model and the TVP-VAR.
Notes: In both cases, the shock is a surprise increase in the interest rate by 25bp. The DSGE model of Alpanda and Zubairy (2017) is calibrated to a "high share" state with an ARM share of 60% and a "low share" state with a share of 10%.

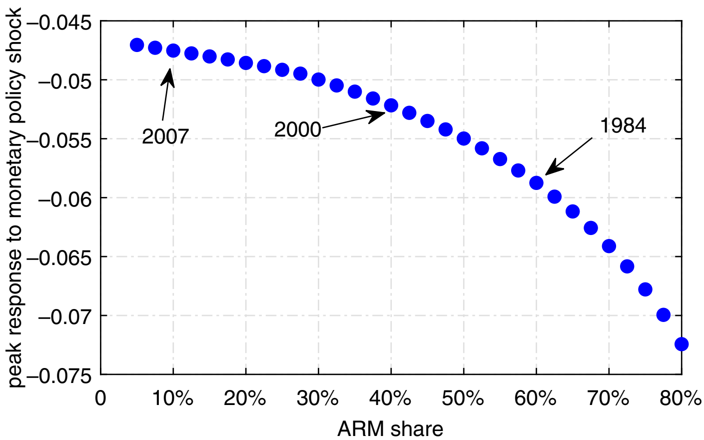

Figure 10. Model-Implied Peak Response of Real Mortgage Debt to Monetary Policy Shock for Alternative ARM Shares.
Notes: Response of real mortgage debt in the DSGE model of Alpanda and Zubairy (2017) for different calibrations of the ARM share. To provide some orientation, we highlight three historical ARM shares.

6. Counterfactual impulse responses

Thus far, our results provide evidence that the transmission of monetary policy shocks may have become weaker over time. Moreover, from the standpoint that our model is supposed to uncover the time-varying structure of the economy, relative cumulative responses provide evidence that periods exist in which reducing debt might be less costly than in others.

One possible driver of a weaker transmission of monetary policy is the structural change in the Fed's policy rule itself. That is, if in the 2000s, the Fed responds systematically different to inflation and unemployment than in the 1970s, the responsiveness of the rest of the economy could change. In this section, we want to rule out that monetary policy itself is driving the change in the transmission mechanism.

We report results for counterfactual experiments that have been widely used (see, for instance, Primiceri, 2005, and Sims and Zha, 2006) and are an informative possibility to establish the role of the Fed in the weaker transmission of monetary policy shocks observed in Section 4. Specifically, we provide counterfactual impulse response functions and counterfactual historical simulations, assuming that the monetary policy rule under Fed chairman Bernanke (February 2006 to January

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

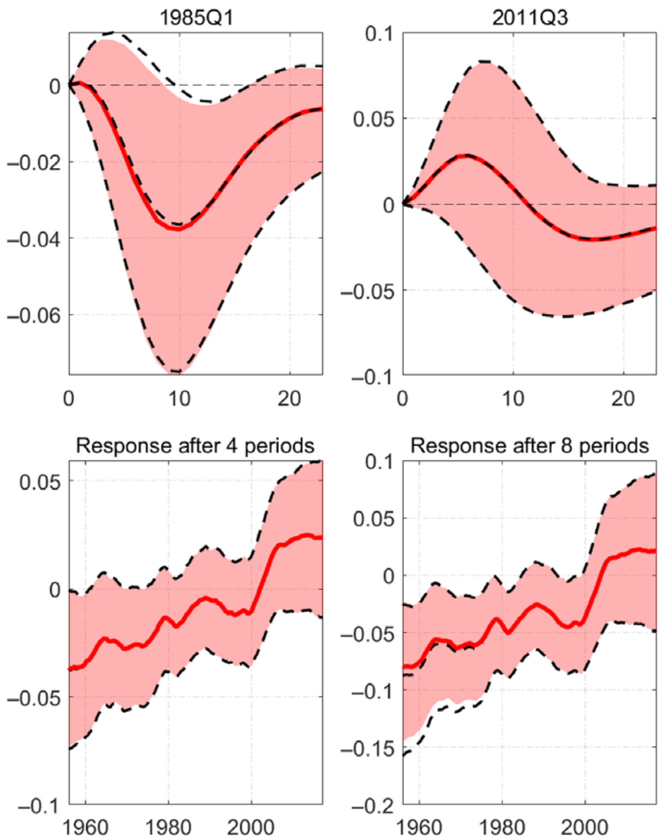

Figure 11. Counterfactual Impulse Responses Under Fixed Policy Rule.
Notes: Mean responses from the baseline model (red-solid) and counterfactual responses (black-dashed) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

2014) prevails over the full estimation sample. We therefore fix the parameters in the interest rate equation of the VAR model to the mean from their posterior during the Bernanke tenure, while all other parameters are allowed to fluctuate.[20]

The counterfactual impulse responses of mortgages are shown in Figure 11, which reports also the baseline results from Figure 2 as a benchmark. For both static impulse responses as well as the responses after four and eight quarters, the impulse responses almost completely overlap with the baseline results. This corroborates the notion that the change in the transmission process does not stem from shifts in the Fed's policy rule.

7. Robustness

The following section examines the sensitivity of our results to a battery of robustness checks. These checks include the estimation of a rolling window VAR model, an alternative identification scheme, an analysis of the non-mortgage component of household debt, the role of the ZLB, and a model with the ratio of mortgage debt to GDP.

7.1 Results of a rolling window estimation

The results in the main section are based on a TVP-VAR model in which both the coefficients and the variance–covariance matrix are time-varying. This provides us with a set of estimated

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

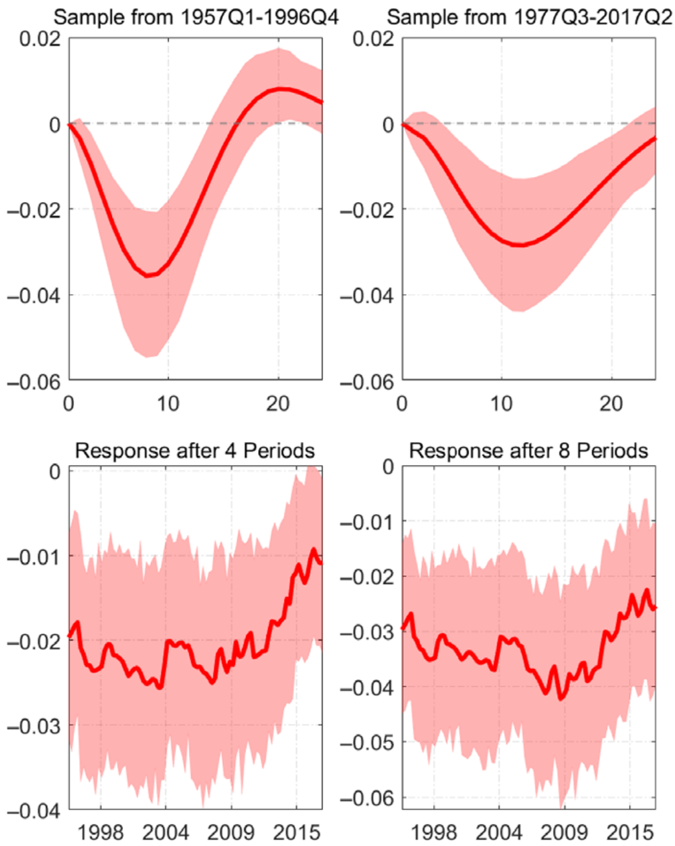

Figure 12. Response of Mortgage Debt from a Rolling Window VAR.
Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock obtained from a rolling window VAR model with two lags.

parameters for each individual period of our estimation sample. In this subsection, we report the results obtained from the estimation of a conventional fixed-parameter VAR over a rolling window.

We estimate the fixed-parameter VAR with the same variables and lag order of two over a fixed estimation window of 40 years, which we move in steps of one quarter over the sample period. After adjusting for lags as well as for the size of the estimation window, we thus obtain 86 sets of estimated parameters. As in the main part of the paper, we identify our shock assuming that the matrix of simultaneous relationships is lower triangular. In order to compare the results based on the same size of the shock, we again assume that an unexpected 25bp monetary policy shock hits the economy. The error bands for the impulse responses are obtained using a standard Monte Carlo simulation with 500 replications.

In order to save space, we only present the most important results here. The lower left panel in Figure 12 shows the response of mortgage debt four periods after the shock over time and the lower right panel shows the same response eight periods after the shock over time. The left panel of the top row shows the impulse response of mortgage debt based on the sample spanning from 1957Q1 to 1996Q4, whereas the right panel shows the response of mortgage debt based on the sample running from 1977Q3 to 2017Q2.

Note that the results for individual periods are not exactly comparable to those of our TVP-VAR as the estimated parameters of the rolling window procedure are average effects over the

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

selected estimation window. As in the main part of the paper, the sensitivity of mortgages to monetary policy shocks decreases significantly over time. That is, our results point to a drop in the responsiveness of mortgage debt to monetary policy. Furthermore, the results show the same pattern as regards the timing of the peak response, that is, the number of periods it takes for the monetary policy shock to reach its maximum impact.[21]

As an additional backup, we repeat the rolling window estimation based on an optimal lag length which we derive from standard information criteria. Based on the Bayesian information criterion, we choose seven lags as an alternative lag length.[22] Our qualitative results remain untouched, as can be seen in the appendix. Although the shape of the impulse responses is not as smooth as before, the trend of decreasing sensitivity of mortgages is more visible under four lags.

As our TVP-VAR runs into difficulties when we estimate the model in trend stationary time series, we further investigate whether our results change when our variables enter the VAR in (log) levels. We use the following variables: (1) real GDP (in logs), (2) the implicit price deflator (in logs), (3) real mortgage debt (in logs), and (4) the effective federal funds rate. We use two lags as suggested by the Bayesian information criterion and the Hannan–Quinn information criterion.

The online appendix summarizes our results for mortgage debt and shows impulse responses for the same samples as above as well as the time-varying effect four and eight quarters after the shock, respectively.[23] Here, too, we see that the sensitivity decreases over time.

Overall, the results from a rolling estimation with different specifications support our findings from the main part of the paper.

7.2 Alternative identifications

Choosing an alternative Choleski ordering can, in principle, affect our results, as we alter the linear combinations of structural shocks which lie behind the reduced-form error terms. The question is whether mortgage debt still shows a negative reaction when we order the short-term interest before mortgage debt. In general, this ordering implies that mortgage debt does not belong to the information set of the central bank. At the same time, however, the contemporaneous reaction of mortgage debt is not restricted as it can react in the same period in which the monetary policy shock occurs.

Figure 13 reports the impulse responses for this ordering for the same selected dates as in Figure 2 as well as the time-varying response four and eight quarters, respectively. Overall, the alternative ordering of variables does not change our results too much. Even when the response of mortgage debt is unrestricted, in almost all periods, we see that the response of mortgage debt shows the same negative sign as in the main part of the paper. In some periods, the contemporaneous response of mortgage debt is not different from zero, suggesting that we get the same qualitative response on impact. We see again that especially toward the end of our sample, the responsiveness of mortgage debt gradually decreases. This trend, however, is less pronounced than in the main part of the paper.

As an alternative identification scheme, we rely on restrictions on the signs of the impulse response functions (see, for instance, Uhlig, 2005, and Arias et al., 2018). We estimate a constant parameter VAR with two lags for the full sample period, using the same variables as in the baseline model. The nature of the restrictions is consistent with a wide consensus in the literature and is backed by standard models of the business cycle. According to our partial identification scheme, a monetary policy shock raises the short rate, leads to a fall in inflation and an increase in unemployment. Finally, the response of mortgages is constrained to zero. All responses are constrained on impact only. These restrictions imply that we cannot observe a price puzzle as was found in the results from our baseline model.

Figure 14 contains the impulse responses. While we cannot learn from this graph about potential time variation, it serves as a consistency check for our baseline results. In particular, we find

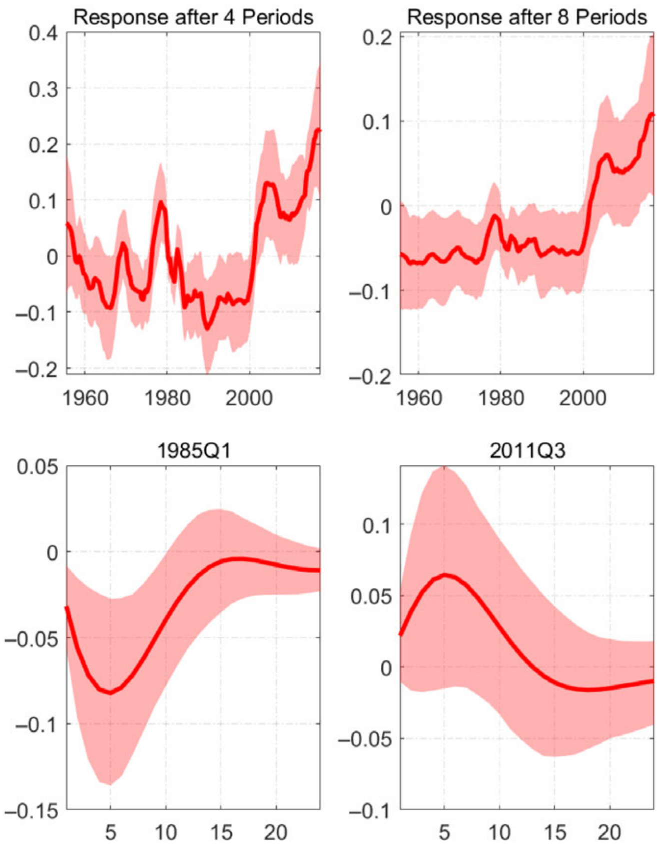

Figure 13. Response of Mortgage Debt for an Alternative Ordering.
Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

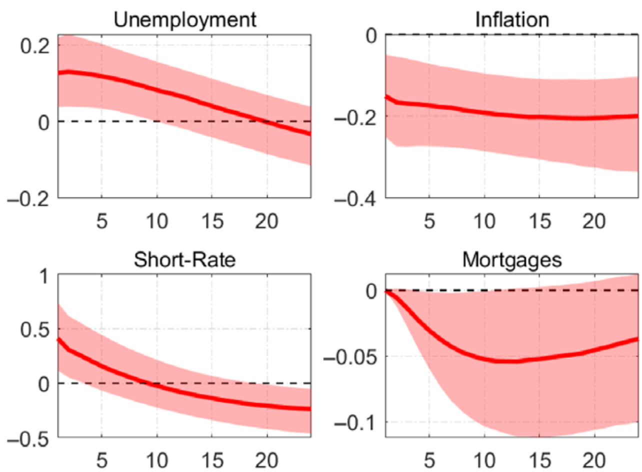

Figure 14. Responses for an Identification Through Sign-Restrictions.
Notes: Mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock.

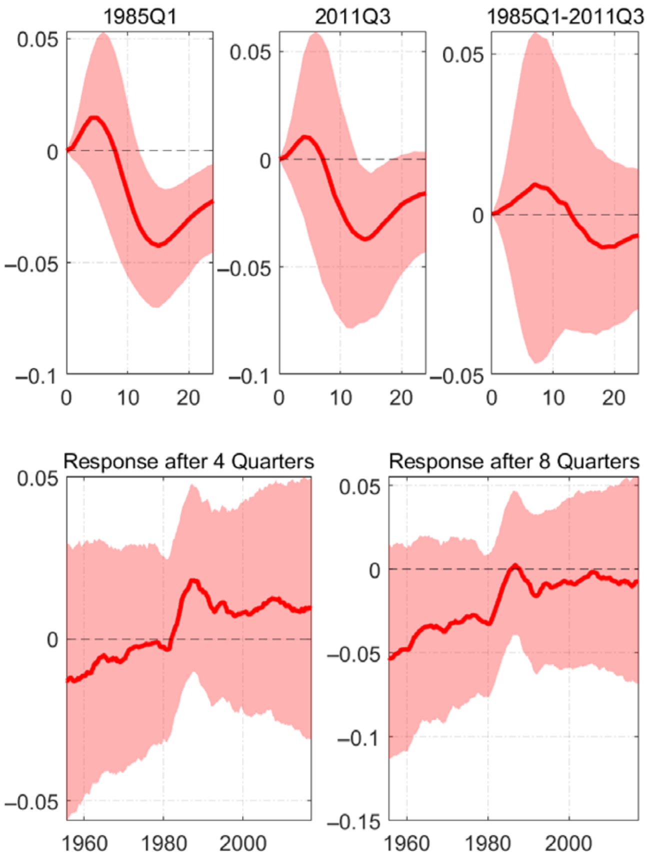

Figure 15. Response of Non-Mortgage Debt.
Notes: Mean responses (red-solid) with 16th and 84th percentiles to a 25bp monetary policy shock.

that mortgages react negatively already one period after the shock. Thus, the direction of the impulse response qualitatively corresponds to the results from the main section. Moreover, the magnitude of the responses roughly equals the sample means of the impulse responses shown earlier.

7.3 Results for non-mortgage debt

We now ask whether the drop in the sensitivity to monetary policy does also hold for non-mortgage household debt. For that purpose, we estimate our baseline model and replace mortgage debt by Baxter and King (1999) filtered non-mortgage household debt.

Figure 15 reports the resulting impulse responses. A policy tightening in 1985Q1 or 2011Q3 leads to a significant contraction of non-mortgage debt. Importantly, both impulse responses are statistically indistinguishable. Hence, we are unable to reject a stable transmission of monetary policy impulses to non-mortgage debt. In addition, the cross-sectional responses shown at the bottom of the figure suggest that the responses after four and eight quarters do not fluctuate over the sample period and remain insignificant. From that, we conclude that the reduced sensitivity of mortgages with respect to monetary policy is indeed a feature of the mortgage market that is not shared by other types of household debt.

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

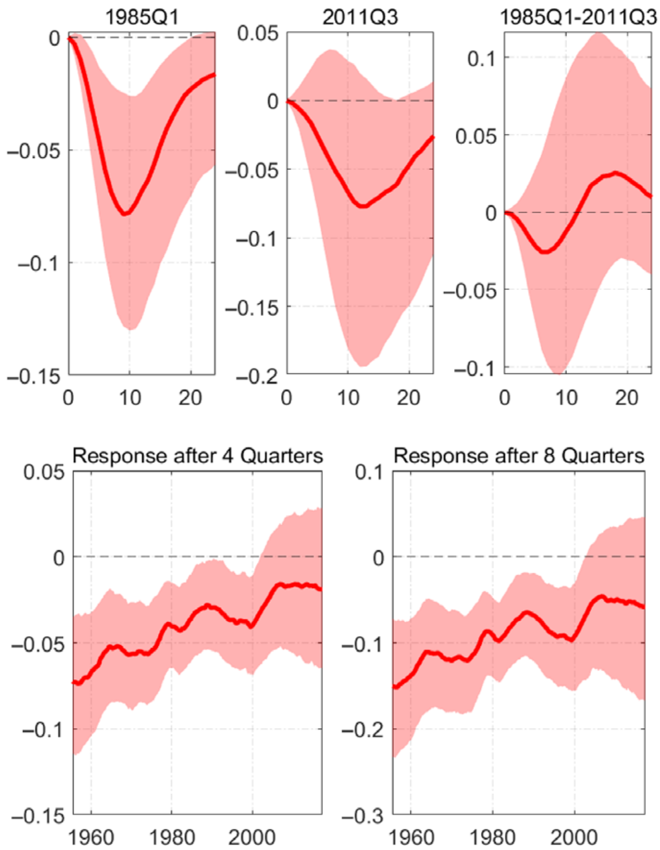

Figure 16. Response of Mortgage Debt: The Role of the ZLB.
Notes: The left panel shows the mean responses (red-solid) with 16^{th} and 84^{th} percentiles to a 25bp monetary policy shock in 1985Q1, whereas the middle panel shows the mean responses to a 25bp monetary policy shock in 2011Q3. The right panel shows the difference between the responses in 1985Q1 and 2011Q3.

7.4 The role of the ZLB

Our sample also covers the period in which the US economy was at the ZLB on nominal interest rates. In the main parts of the paper, we show results for a VAR model that includes the federal funds rate for most of the sample period but uses a shadow short rate for the ZLB episode. We now drop the shadow short rate and use the federal funds rate for the full estimation sample. Hence, our model includes a period when the short rate was constrained by zero.

The resulting impulse responses of mortgage debt are shown in Figure 16. We find a pattern that is very close to our baseline results. Therefore, we can conclude that the ZLB or the use of the shadow rate, respectively, does not affect our findings.

7.5 A model with mortgage debt to GDP

The analysis thus far rests on the cyclical component of real mortgage debt, which we obtain from band-pass filtering the original time series. Alternatively, one could express real mortgage debt relative to real GDP. As our estimation framework relies on stationary time series, however, we need to include this series in first differences.

https://doi.org/10.1017/S136510052100050X Published online by Cambridge University Press

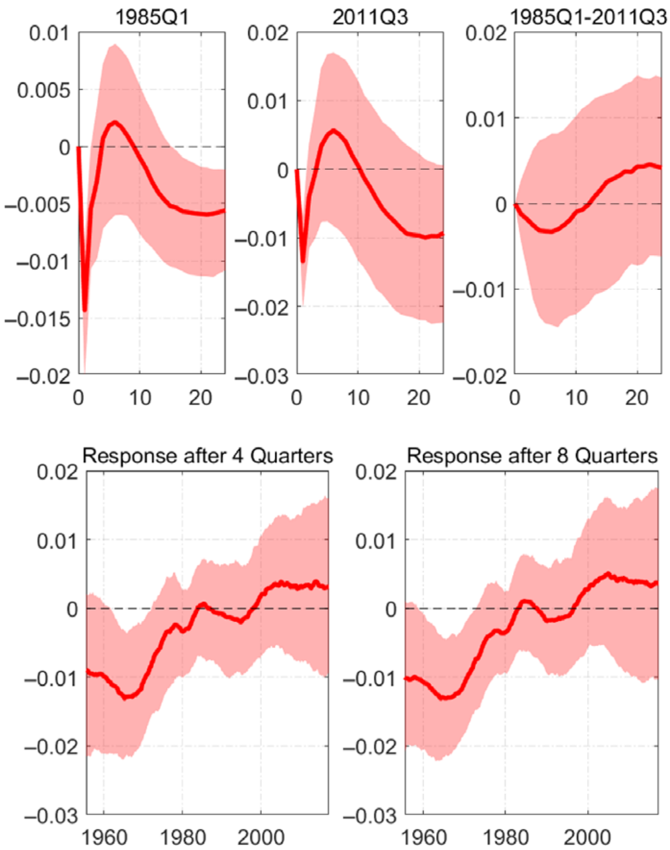

Figure 17. Response of Mortgage Debt relative to GDP.
Notes: Mean responses (red-solid) with 16th and 84th percentiles to a 25bp monetary policy shock.

Figure 17 depicts the key results, that is, the responses of mortgage debt to GDP at two points in time as well as the evolution of the cross-sectional responses over time. The similarity with our baseline findings is striking: again, we find that mortgage debt fell stronger after a policy shock in 1985 compared to the same shock in 2011. Furthermore, the cross section of the impulse responses document a decline in the effectiveness of monetary policy which is qualitatively similar to the results from our baseline model.

8. Conclusions

In this paper, we studied the role of monetary policy for the dynamics of US mortgage debt, the largest and most important component of overall household debt. In the aftermath of the recent financial crisis, which originated in the US housing market, the mortgage market received much attention.

The main tool of our analysis, a time-varying VAR model with stochastic volatility, allowed us to study the sensitivity of mortgage debt to monetary policy over time. We find that since the 1960s, the impact of monetary policy on mortgage debt has declined. A policy shock in 2011 has a much smaller effect on mortgage debt than an identically sized shock occurring in 1985. This finding, which is new to the literature, is robust to variations of the model and the estimation approach and is not observed for non-mortgage debt.

We also estimate a DSGE model for the US economy in order to replicate our empirical findings. The ARM share, a key parameter in the determination of the model-based impulse responses, is shown to have declined strongly since the early 1980s. Once we calibrate the model to alternative realizations of the ARM share, we are able to replicate the decline in the response of debt to monetary policy quantitatively. To the extent the ARM share could be taken as given, this offers a consistent explanation for our results.

These findings have several implications for monetary policy and the mortgage market. First, our results suggest that, nowadays, monetary policy is a blunt and inefficient tool to engineer a reduction of household debt. The decline in the sensitivity to monetary policy implies that a large policy adjustment is needed in order to have a sizable effect on mortgage debt. This, however, would cause a deep recession. Hence, our results speak against using monetary policy as an instrument to prevent the build-up of household debt.

A second interpretation of our results addresses the role of the Fed in the run-up to the recent financial crisis. It is often claimed that the Fed contributed to inflating house prices by keeping the federal funds target rate too low for too long. Our results put this claim into perspective. If the sensitivity of mortgage debt to monetary policy in the mid-2000s is low, which is our main result, even persistently low levels of the federal funds rate should contribute little to the rise in mortgage debt before the crisis. Likewise, tightening monetary conditions, as the Fed did after June 2004, should translate into a small decrease in mortgage debt.

Our results fit the "mortgage rate conundrum" diagnosed by Justiniano et al. (2017). These authors argue that the empirical link between mortgage rates and longer-term interest rates broke. Hence, there seem to be strong structural changes in the mortgage market and its link to monetary policy.

Acknowledgements. We thank the anonymous referees, Mathias Klein, Jouchi Nakajima, and Sarah Zubairy for important input. Participants of the 2nd Workshop on Financial Econometrics and Empirical Modeling of Financial Markets, Kiel Institute for the World Economy, the WEAI 93rd annual conference in Vancouver, the 33rd annual conference of the EEA in Cologne, the 24th CEF conference in Milan and of seminars at University Giessen provided important feedback. We thank the Fritz Thyssen Stiftung for financial support (grant no. 10.16.2.004WW).

Supplementary materials. To view supplementary material for this article, please visit http://dx.doi.org/10.1017/10.1017/S136510052100050X.

Notes

1 See Jordà et al. (2015) for historical evidence of the link between loose monetary policy and real estate lending booms.
2 Huber and Punzi (2020) also use a TVP-VAR model to study the transmission of monetary to the housing market at the zero lower bound.
3 Eickmeier and Hofmann (2013) use a factor-augmented VAR model to study the monetary policy transmission to the housing market. They show that monetary policy contributed to the housing precrisis housing boom.
4 Harding and Klein (2019) also study the nonlinear impact of monetary policy on the US housing market. They show that monetary policy is effective in periods of deleveraging but not during the build-up of household leverage.
5 The appendix provides additional details on the sources of the variables.
6 Primiceri (2005) uses the yield on 3-month Treasury bills as opposed to the federal fund as in this paper.
7 Note that the lower-triangular specification (2) of $\mathbf{A}$ (and thus $\mathbf{A}^{-1}$) is widely used and enables us to easily identify structural shocks (e.g., monetary policy shocks) by recursive ordering, although the examination of implications for the economic structure may require more complicated identification schemes.
8 It turns out that even if we do not impose these restrictions, our results remain indistinguishable.
9 However, while the marginal likelihood might be one natural candidate to find the optimal lag length within a TVP-VAR, several articles criticized that the harmonic mean method (as was typically used in the literature) could be biased. Therefore, a rolling window estimation can provide an alternative indication of the optimal lag length, because the calculation of the marginal (or log) likelihood within time-invariant VAR models is tractable and well understood. Based on an estimation window of size 160, for each estimation, we calculate and store standard information criteria over time. It turns out that this procedure suggests a lag length of four. Although we estimate our baseline model with two lags in order not to overload the model, we also estimate our TVP-VAR with a lag length of $s = 4$. Besides a significantly increased estimation time, our main results remain qualitatively mostly similar, as can be seen in the appendix.

10 See Nakajima (2011) and Primiceri (2005) for detailed discussions.
11 Note that the problem of potential spurious movements is negligible as we choose rather tight priors for the covariance matrix of the disturbance in the random walk process.
12 Accordingly, $\mathbf{B} = \{\mathbf{B}_{s+1}, ..., \mathbf{B}_T\}$, $\mathbf{a} = \{\mathbf{a}_{s+1}, ..., \mathbf{a}_T\}$ and $\mathbf{h} = \{\mathbf{h}_{s+1}, ..., \mathbf{h}_T\}$.
13 For a useful exposition of the MCMC algorithm, see Nakajima (2011).
14 We therewith follow Nakajima (2011) by separate priors for the i^{th} diagonals of the covariance matrices in the transition equations. Primiceri (2005), on the other hand, specifies priors for $\mathbf{\Sigma}_B$, $\mathbf{\Sigma}_a$, and $\mathbf{\Sigma}_h$ based on an inverse-Wishart distribution. Note that the Wishart distribution is nothing but a multivariate generalization of the gamma distribution. In this case, drawn samples are positive-definite matrices rather than positive real numbers.
15 In the online appendix, we also show that the TVP-VAR model delivers a superior fit to the data compared to a constant parameter VAR model.
16 The calculation of the peak timing is based on the posterior mean. In a previous version, we showed the median of the peak timing.
17 Moench et al. (2010) discuss the reduction in the ARM share in the USA and explain it in terms of financial innovations such as an increase in securitization and a shifting term structure of interest rates.
18 In the next subsection, we simulate a structural model to corroborate our findings.
19 In the appendix, we explain in detail how we account for different values for the ARM share.
20 Figure (1) in the appendix shows the corresponding counterfactual historical simulation for the exercise described above. Here, too, we see that a possible shift in the policy rule has hardly contributed to a change in the historical development of our endogenous variables. Only the unemployment rate would have been somewhat lower with the "Bernanke Rule" in the 1980s. The difference is not significant, however, as the observed data lie within the 16^{th} and 84^{th} percentiles of the counterfactual simulation.
21 The impulse responses for the other variables are available on request. However, it stands out that the price puzzle disappears as the impulse response of the inflation rate now shows the expected negative sign much earlier than before.
22 We allow for a maximum lag length of eight lags.
23 The results are similar when we use the unemployment rate instead of (log) real GDP.

References

Alpanda, S. and S. Zubairy (2017) Addressing household indebtedness: monetary, fiscal or macroprudential policy? *European Economic Review* 92, 47–73.

Alpanda, S. and S. Zubairy (2019) Household debt overhang and transmission of monetary policy. *Journal of Money, Credit and Banking* 51(5), 1265–1307.

Arias, J. E., J. F. Rubio-Ramírez and D. F. Waggoner (2018) Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. *Econometrica* 86(2), 685–720.

Baxter, M. and R. G. King (1999) Measuring business cycles: Approximate band-pass filters for economic time series. *Review of Economics and Statistics* 81(4), 575–593.

Ben Zeev, N. (2016) Household Debt, Adjustable-Rate Mortgages, and the Shock-Absorbing Capacity of Monetary Policy. *unpublished.*

Boivin, J. (2006) Has US monetary policy changed? Evidence from drifting coefficients and real-time data. *Journal of Money, Credit and Banking* 38(5), 1149–1174.

Calza, A., T. Monacelli, and L. Stracca (2013) Housing finance and monetary policy. *Journal of the European Economic Association*, 11(suppl. 1), 101–122.

Canova, F. and L. Gambetti (2009) Structural changes in the US economy: Is there a role for monetary policy? *Journal of Economic Dynamics and Control* 33(2), 477–490.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999) Monetary policy shocks: What have we learned and to what end? In: J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, vol. 1, Part A, 65–148.

Cogley, T. and T. J. Sargent (2005) Drifts and volatilities: Monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8(2), 262–302.

Ehrmann, M. and M. Ziegelmeyer (2017) Mortgage choice in the euro area: Macroeconomic determinants and the effect of monetary policy on debt burdens. *Journal of Money, Credit and Banking* 49(2-3), 469–494.

Eickmeier, S. and B. Hofmann (2013) Monetary policy, housing booms, and financial (im)balances. *Macroeconomic Dynamics* 17, 830–860.

Garriga, C., F. E. Kydland and R. Šustek (2017) Mortgages and monetary policy. *The Review of Financial Studies* 30(10), 3337–3375.

Geweke, J. (1992) Evaluating the accuracy of sampling-based approaches to the calculations of posterior moments. *Bayesian Statistics* 4, 641–649.

Harding, M. and M. Klein (2019) Monetary Policy and Household Deleveraging. *unpublished*.

Huber, F. and M. T. Punzi (2020) International housing markets, unconventional monetary policy, and the zero lower bound. *Macroeconomic Dynamics* 24, 774–806.

Iacoviello, M. (2005) House prices, borrowing constraints, and monetary policy in the business cycle. *American Economic Review* 95(3), 739–764.

Jordà, Ò., M. Schularick and A. M. Taylor (2015) Betting the house. *Journal of International Economics* 96, S2–S18.

Jordà, Ò., M. Schularick and A. M. Taylor (2016) The great mortgaging: Housing finance, crises and business cycles. *Economic Policy* 31(85), 107–152.

Justiniano, A., G. E. Primiceri and A. Tambalotti (2017) The Mortgage Rate Conundrum. *unpublished*.

Koop, G. and D. Korobilis (2010) *Bayesian Multivariate Time Series Methods for Empirical Macroeconomics*, Foundations and Trends in Econometrics.

Mian, A., A. Sufi and E. Verner (2017) Household debt and business cycles worldwide. *The Quarterly Journal of Economics* 132(4), 1755–1817.

Mishkin, F. S. (2009) Globalization, macroeconomic performance, and monetary policy. *Journal of Money, Credit and Banking* 41, 187–196.

Moench, E., J. I. Vickery and D. Aragon (2010) Why is the market share of adjustable-rate mortgages so low? *Current Issues in Economics and Finance* 16 (8).

Nakajima, J. (2011) Time-varying parameter VAR model with stochastic volatility: An overview of methodology and empirical applications. *Monetary and Economic Studies* 29, 107–142.

Owyang, M. T. and T. Sekhposyan (2012) Okun's law over the business cycle: Was the great recession all that different? *Federal Reserve Bank of St. Louis Review* 94(September/October 2012).

Paul, P. (2020) The time-varying effect of monetary policy on asset prices. *Review of Economics and Statistics* 102(4), 690–704.

Primiceri, G. E. (2005) Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* 72(3), 821–852.

Raftery, A. and S. Lewis (1992) How many iterations in the Gibbs sampler? *Bayesian Statistics* 4, 765–776.

Sims, C. A. and T. Zha (2006) Were there regime switches in U.S. monetary policy? *American Economic Review* 96(1), 54–81.

Uhlig, H. (2005) What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics* 52(2), 381–419.

Wooldridge, J. M. (2016) *Introductory Econometrics: A Modern Approach*. Nelson Education.

Wu, J. C. and F. D. Xia (2016) Measuring the macroeconomic impact of monetary policy at the zero lower bound. *Journal of Money, Credit and Banking* 48(2–3), 253–291.

Cite this article: Finck D, Schmidt J and Tillmann P (2023). "Mortgage debt and time-varying monetary policy transmission", *Macroeconomic Dynamics* **27**, 506–531. https://doi.org/10.1017/S136510052100050X