Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

Inframarginal Borrowers and the Mortgage Payment Channel of Monetary Policy

Daniel Ringo

2024-069

Please cite this paper as:
Ringo, Daniel (2024). "Inframarginal Borrowers and the Mortgage Payment Channel of Monetary Policy," Finance and Economics Discussion Series 2024-069. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2024.069.

Inframarginal Borrowers and the Mortgage Payment Channel of Monetary Policy

Daniel Ringo*
Federal Reserve Board
July 30, 2024

Abstract

Despite the widespread use of fixed-rate mortgages in the United States, I show that monetary policy is effectively passed through to aggregate outstanding mortgage debt service. Using credit bureau, lender, and servicer data on mortgage payments and originations and exogenous monetary policy shocks, I estimate a mortgage rate semi-elasticity of payments over 10. Inframarginal borrowers—households whose choice to buy a home or refinance does not depend on the particular monetary policy decision under consideration—are the most important conduit, explaining over half of the pass-through. Consistently large flows of inframarginal borrowing relative to the stock of outstanding debt account for the strength of this channel. Households with adjustable-rate mortgages and marginal refinancers, the focus of much of the literature on monetary policy's effect on mortgage borrowers, each explain about 20 percent of the pass-through. I show the mortgage payment channel induces a lag in the operation of policy, as the cumulative effects on debt service build over time in response to persistent shocks to longer-term rates. Estimated magnitudes suggest that mortgage payments are a primary channel by which monetary policy affects consumption.

Keywords: Monetary policy; interest rates; refinancing channel; debt service

JEL codes: G21; E43; E52.

*MS 93, Board of Governors of the Federal Reserve System, Washington D.C. 20551 daniel.r.ringo@frb.gov. Thanks to Neil Bhutta, John Driscoll, and seminar participants at the Federal Reserve Board and Ghent University Empirical Macroeconomics Workshop for helpful comments. The analysis and conclusions set forth are those of the author and do not indicate concurrence by other members of the research staff or the Board of Governors.

1 Introduction

Monetary policy can affect households' consumption via their debt service payments, increasing their disposable incomes net of required payments as interest rates decrease and vice versa. Residential mortgage debt represents the outright majority of all household debt and is the largest source of debt service obligations, and so should be a potent channel by which monetary policy can influence consumer demand. In the United States, however, the popularity of fixed-rate mortgages presents an obstacle to the pass-through of interest rate changes to payments. When rates rise, debtor households will hang on to their old rates and be unaffected. When rates fall, only a fraction of households for whom refinancing would be profitable exercise their option to do so. The frictions inherent in this process have motivated a growing body of literature analyzing the refinancing channel of monetary policy.[1] In an international context, it is commonly argued that the prevalance of fixed-rate mortgages dampens the effectiveness of monetary policy in countries like the U.S.[2]

In this paper, I show that the pass-through of monetary policy shocks to mortgage payments in the U.S. is actually quite robust due to the large volume of borrowing by inframarginal borrowers. These are households whose decision to take out a new loan does not depend on the particular monetary policy action under consideration, but whose payment for that new loan still depends on the stance of policy. Importantly, inframarginal borrowers include households who are borrowing for the purposes of home purchase, refinancing, or any other ends. While only a small percentage of households engage in any new mortgage borrowing in a given year, in dollar terms annual new borrowing represents approximately one quarter of the outstanding stock of household mortgage debt. This is because new loans tend to be larger than the unpaid principal balance of existing loans, and explains the strength of the mortgage payment channel. Historically, a large percentage of mortgage debt has carried a recently determined interest rate.

The effect of a typical monetary policy shock on the monthly payments of an inframarginal borrower is small relative to that experienced by a marginal refinancer—a household

[1] See Amromin, Bhutta and Keys (2020) for a helpful summary.

[2] See, e.g., Bernanke (2007), Calza, Monacelli and Stracca (2013), Rubio (2011), or Kim and Lim (2020).

that would be induced into or out of refinancing their current mortgage by a small change in rates. However, inframarginal borrowers are so much more numerous than marginal borrowers that the inframarginal channel is more than twice as important for the pass-through of rate shocks to aggregate payments. I find that over the last 25 years, a monetary policy shock which cuts mortgage rates by only 7 basis points would reduce aggregate annual mortgage payments by approximately $9 billion in 2023 dollars. Inframarginal borrowers would account for about $5 billion of this reduction, and marginal refinancers $2 billion. An additional $2 billion would go to the minority of borrowers with adjustable-rate mortgages (ARMs), although this latter channel has diminished in recent years. Marginal home purchase borrowers and the intensive margin of borrowing (i.e., the amount borrowed conditional on getting a loan) are not significant channels of pass-through to payments in the first few years post-shock.

The stance of monetary policy and the level of mortgage rates are endogenous to the state of the economy. To consistently estimate the effects of policy decisions, I therefore exploit variation from the exogenous, unanticipated components of rate responses to Federal Open Market Committee (FOMC) announcements. I begin by estimating the effects of monetary policy shocks on mortgage interest rates. It is important here to recognize that mortgage interest rates are tied to yields on relatively longer-maturity debt (as a typical mortgage will endure for years before prepaying) and the Federal Reserve has tools at its disposal that affect different portions of the yield curve differently. The Fed can directly control very short-term rates by setting the overnight federal funds rate. Longer-term rates depend more on investors' beliefs about the future path of monetary policy decisions, and the Fed can attempt to influence these beliefs via forward guidance and asset purchases. I therefore use the series of monetary policy shocks estimated by Swanson (2021), who separately identified shocks to the federal funds rate, forward guidance, and large-scale asset purchases (LSAP) contained in FOMC announcements.[3] Using the local projections estimator of Jordà (2005),

[3]Lewis (2023) also estimates shock series that disentangle these components using different identifying assumptions than Swanson (2021). In addition, Lewis (2023) attempts to isolate a fourth component, the Fed information effect. In Appendix Section A.1 I show my results are robust to using these alternative shocks series.

I estimate that a one standard deviation forward guidance shock raises mortgage rates by an average of 7.4 basis points over the three years post-shock.[4] Shocks to the federal funds rate do not show any demonstrable effect on mortgage rates, however. The LSAP shock is estimated to have a strong initial effect on mortgage rates, but this quickly fades.

Next, I estimate how the average payments of mortgage borrowers are affected by monetary policy shocks empirically. I construct a time series of the average mortgage payments of all outstanding mortgage borrowers using nationally representative credit bureau data from 1999 through 2023. Again using the local projections estimator, I estimate that a one standard deviation forward guidance shock increased the average payments of borrowers by 0.89 percent after three years, or approximately $9 billion per year. Putting the estimates together, each basis point increase in mortgage rates is estimated to increase annual payments by over $1 billion; in other words, the interest rate semi-elasticity of mortgage payments is about 10.

The estimated effects on mortgage debt service are large enough to explain much of the consumption response to monetary policy shocks. Cloyne, Ferreira and Surico (2020) estimate that in the U.S. a 25 basis point generic monetary policy shock would decrease durable and non-durable consumption by approximately 0.9 and 0.15 percent, respectively, three years post-shock. With receent annual personal consumption expenditures of approximately $2 trillion and $17 trillion in these categories, this implies a total reduction of about $45 billion per year. My estimates suggest that a 25 basis point shock to mortgage rates would increase annual mortgage payments by about $30 billion by year three. If the marginal propensity to consume out of a persistent change in disposable income is high, most of the effect of monetary policy on consumption can therefore be explained by the direct effect of the mortgage payment channel.[5] Any multiplier effects from the shock to mortgage borrowers'

[4]Prior studies, such as Hamilton (2008), Gorea, Kryvtsov and Kudlyak (2022), and Ringo (2024), have shown the effect of forward guidance on mortgage rates in the short-term, but this study is the first I am aware of to demonstrate that the effect persists for years.

[5]While estimates of the marginal propensity to consume vary in the literature, most researchers find that it is high for permanent shocks to income. See Jappelli and Pistaferri (2010) for a review of the literature. Changes in mortgage payments represent a transfer between borrowers and creditors, so there are potentially offsetting consumption effects. However, most mortgage debt in the US (often packaged into mortgage backed securities) is held by commercial banks, the Federal Reserve, or institutional investors, all of whom presumably have a low marginal propensity to consume.

consumption would further increase the importance of the mortgage payment channel. The strength of the mortgage payment channel also provides an explanation for monetary policy's effect on consumption despite micro evidence suggesting a low intertemporal elasticity of substitution (see, e.g., Best et al., 2020).

The total estimated change in payments includes the effects on inframarginal borrowers, marginal borrowers, and existing ARMs (whose mortgagors do not need to engage in any new originations to be affected). Additionally, households engaging in new borrowing may offset some of the effect of monetary policy on payments by making an intensive margin adjustment, i.e. reducing the size of their loans in the face of higher rates. To identify the relative importance of these channels, I estimate their magnitudes separately.

Inframarginal borrowers are simple to model because, by definition, their borrowing decisions are exogenous to the rate shock. The magnitude of the inframarginal borrower channel can therefore be straightforwardly estimated via a simple back-of-the-envelope calculation. Using the average origination volume and mortgage interest rate over the past 25 years, I calculate that cutting rates by 7.4 basis points for three years should reduce aggregate mortgage debt service by approximately $5.2 billion per annum. The effect builds linearly over time so long as the lower rate is maintained, as each month a new cohort of inframarginal borrowers are affected while cohorts that previously got their loan continue to pay the lower rate. Comparing this simple calculation to the total estimated change in payments, it seems that inframarginal borrowers are reponsible for just over half of the pass-through of monetary policy.

In addition to the increased or decreased payments faced by these inframarginal borrowers, a rate shock can pull in or exclude marginal borrowers. The number of such borrowers can be estimated by the responsiveness of origination volumes to rate shocks—the difference between this estimated number and the total number of originations is the number of inframarginal borrowers. Marginal borrowers are vastly outnumbered by inframarginal borrowers for rate shocks of a typical historical magnitude, but the effect on their payments per-borrower may be much larger. This is because, for marginal borrowers, the counterfactual payment is their current existing payment (with its old rate potentially very different from current rates), not the payment they would get originating the same loan with a marginally different rate (as it would be for inframarginal borrowers).[6]

To determine the magnitude of the marginal refinance channel, I estimate the effect of the same monetary policy shock on the number of weekly refinance originations over the 1999-2023 period. Again using a local projections estimator, I find the shock decreases refinancing volumes by 3.5 percent over the three years post-shock, suggesting approximately $2 billion of the total increase in payments comes from marginal refinancers. For the marginal home purchase channel, I run the same estimator on the number of weekly home purchase originations. I do not find effects on home purchase lending volumes that are either statistically significant or economically consequential for average payment amounts.

To determine the magnitude of the ARM channel, I use a monthly panel of servicing records of individual ARMs from 2000 through 2023. The local projections estimator indicates that the same monetary policy shock increases ARM payments by approximately 1.2 percent, equivalent to $2 billion in additional payments per year.

The back-of-the-envelope calculation of the inframarginal borrower channel ignores the possibility that households adjust on other margins (i.e., loan size). To test this implicit assumption, I next estimate the magnitudes of the intensive margin channel. To do so, I run the local projections estimator on a weekly series of the average loan size of new originations. I find effects that are close to zero for the first few years after the shock. The assumption of no intensive-margin adjustment used in the back-of-the-envelope calculation of the inframarginal channel appears benign.

Together, the estimated magnitudes of the inframarginal borrower, marginal refinancer, and ARM channels add up to be close to the estimated net effect on aggregate payments. The estimates of channel-specific magnitudes use outcome measures wholely independent of

[6]For this reason, monetary tightenings can increase payments via the marginal refinancing channel (relative to the counterfactual of no tightening) nearly symmetrically to the manner in which monetary easings lower payments. If a small tightening shock dissaudes a borrower from refinacing, they forego a relatively large reduction in payment. For example, suppose a borrower would choose to refinance only if rates fall to one full percentage point below that borrower's existing rate. A 5 basis point shock that moved rates from 0.95 percent to 1.00 percent below their rate would induce the borrower to refinance and decrease their existing rate by 1 percentage point. Symmetrically, a 5 basis point tightening shock from 1.00 to 0.95 below their rate would dissuade that same borrower from refinancing, increasing their rate by 1 percentage point relative to the counterfactual of no shock.

the time-series measure of payments used to estimate the overall effect. The fact that the result from the "top-down" approach to estimating the magnitude of the mortgage payment channel matches that from the "bottom-up" approach was not externally imposed, and confirms the finding is robust. Interestingly, all three of these channels naturally introduce a lag between the monetary policy decision and the maximum effect on payments. For the inframarginal borrowers and marginal refinancer channels, this occurs because the change in payments accumulates slowly as each month a new set of loans is originated subject to, or contingent on, the policy-affected rate. For the ARM channel, a lag occurs because ARM rates only adjust at periodic intervals (often one year), or even less frequently if the loan is still in an initial fixed-rate period. Milton Friedman's "long and variable" lags have a natural explanation, at least in recent decades, as a consequence of the features of the mortgage market on which monetary policy operates.

Finally, I turn to a consideration of the state dependence of monetary policy's effectiveness through the inframarginal borrower channel. The total effect of a rate shock on payments varies proportionately with the volume of mortgage borrowing Americans will engage in over the subsequent three years. This volume has varied considerably over the past quarter century, with three-year peaks of over $13 trillion (in real 2023 dollars) reached during the refinancing booms of 2003 and the Covid pandemic, while the average three-year volume was closer to $9 trillion. Three-year volumes never fell below $6 trillion, however, as less-volatile home purchase borrowing supports the floor of total borrowing even when refinancing enters a fallow period. These flows are always quite large relative to the stock of households' outstanding mortgage debt (approximately $13 trillion[7]), suggesting that the fixed-rate nature of most mortgage debt only partially impedes the pass-through of rate changes to household disposable income. In addition to ARMS and fixed-rate mortgages exhibiting similar lags in the pass-through of monetary policy shocks to payments, the per-borrower effect sizes are comparable. Due to the large flows of new originations, I estimate the payments of the average mortgage borrower are about 75 percent as responsive to a forward guidance shock as are those of an ARM borrower.

[7]Source: Federal Reserve Board (2023)

My paper makes several contributions. First, this is the only paper I am aware of to demonstrate using quasi-experimental methods the existence of a strong mortgage payment channel of monetary policy in the U.S. It is also the first I am aware of in this literature to attempt to empirically distinguish and quantify the magnitudes of the various sub-channels by which monetary policy is passed through to mortgage payments. My finding that the inframarginal channel is substantially the largest should shape how we think about the transmission of monetary policy decisions to household spending. Furthermore, to my knowledge this is the first study to demonstrate that the effect of central bank forward guidance on mortgage rates persists for several years. This persistence is key to the overall strength of the mortgage payment channel, as the effects on payments through all of the sub-channels I identify grow cumulatively the longer rates are held above or below their counterfactual level.

1.1 Related Literature

A large body of literature has investigated the means by which monetary policy affects the mortgage market, and how these changes in turn affect consumption. Studies that estimate how refinancing to a lower rate increases household consumption include Di Maggio et al. (2017), Di Maggio, Kermani and Palmer (2020), Abel and Fuster (2021), and Agarwal et al. (2023). Other authors have investigated the heterogeneous effects of monetary policy by differential household propensities to refinance, including Beraja et al. (2019), Wong (2019), and Cumming and Hubert (2021).

Among these, a paper that is particularly related to the present study is Cloyne, Ferreira and Surico (2020). The authors study how monetary policy shocks affect the consumption of households by their ownership and mortgage status in the U.S. and the U.K. They also estimate how shocks affect average mortgage payments of borrowers with outstanding loans, as I do in this paper. However, Cloyne, Ferreira and Surico (2020) find much smaller effects of monetary policy shocks on payments than I do. They find a monetary policy shock that cuts the federal funds rate by 25 basis points lowers payments by only $3.60 per month for the average borrower.[8] In contrast, I estimate that a mere 7 basis point mortgage rate cut induced by forward guidance would lower average payments by about $18 per month after three years.[9]

As I show in Appendix Section A.2, this difference in our findings is likely due to Cloyne, Ferreira and Surico (2020)'s use of survey responses in the Consumer Expenditure Survey (CEX) to construct their time series of mortgage payments. In comparison to the credit bureau data I employ, the CEX has a much smaller, frequently rotating sample and relies on individual borrowers rather than loan servicers accurately reporting their financial data. I show that the CEX-based payments series is much noisier than the one I construct from the CCP. When I run my estimator on the CEX data, I do not find a significant effect of monetary policy shocks on reported mortgage payments. However, a null net effect of rates on payments is hard to square with the magnitudes of the individual channels I estimate in Section 5, and this finding is most plausibly attributed to small-sample limitations of the CEX.

Furthermore, my results help rationalize a surprising tension in Cloyne, Ferreira and Surico (2020)'s findings: despite a limited estimate of the effect of monetary policy on mortgage payments in the U.S., they find that the consumption response of mortgage borrowers to monetary shocks is much larger than that of renters or outright owners. This latter result is more intuitive in light of the large mortgage payment channel my findings demonstrate.

A further strand of the literature has studied how the embedded-option nature of fixed-rate mortgages implies the refinancing channel of monetary policy is heavily state dependent. Because borrowers are only likely to refinance when their loan is "in the money" (i.e., their current rate is well above the prevailing market rate), monetary policy is less able to affect borrower's payments when rates have risen from long-term lows. Berger et al. (2021) and Eichenbaum, Rebelo and Wong (2022) model how this complication can hamper the ability

[8]Cloyne, Ferreira and Surico (2020) report an effect size of $1.20 in 1984 dollars, and I adjust for the 300 percent inflation between 1984 and 2023.

[9]This estimate is calculated by applying my estimate of a 0.89 percent payment response to a one standard deviation forward guidance shock to an average borrower's monthly payment of $2,000, calculated as described in Section 2. Effects from the inframarginal channel alone would increase payments by approximately $10 per month per borrower.

of rate cuts to stimulate consumption. My findings support their main results, although I point out that the overall inframarginal borrowing channel is not as state dependent as the refinancing channel alone. Even when very few borrowers are in the money to refinance, total borrowing volumes are still supported by relatively rate-inelastic home buying demand and monetary policy continues to have new loans to operate on.

Marginal cash-out refinancing is another related channel by which monetary policy could influence spending, as Bhutta and Keys (2016) find equity extraction is quite responsive to interest rates. However, Anenberg, Scharlemann and van Straelen (2023) show that, at least since the end of the housing boom of the early 2000s, much or all of the marginal equity extracted due to mortgage rate shocks is a substitute for other forms of consumer borrowing, and has little net effect on consumption. In addition to refinancing, studies that estimate the effect of interest rates and monetary policy on home purchase borrowing include Bhutta and Ringo (2021) and Ringo (2024) on the extensive margin, and DeFusco and Paciorek (2017) on the intensive margin.

In an international context, studies of the debt-service channel of monetary policy include Hofmann and Peersman (2017) and Flodén et al. (2021).

2 Data

Data for this project come from several sources. To construct a time series of average monthly payments for households with a mortgage, I use the Federal Reserve Bank of New York/Equifax Consumer Credit Panel (CCP). The CCP consists of a quarterly panel of credit bureau records of a randomly selected, anonymized 5 percent sample of all adults in the U.S. with a credit record. Among other information, loan servicers report their borrowers' scheduled debt service payments for each trade line to the credit bureaus. Over the sample period 1999 through 2023, the CCP includes data on 2.5 million to 3.6 million primary sample individuals with a positive mortgage payment

For each quarter from 1999Q1 through 2023Q4, I take the average of total reported monthly payments on first lien mortgages per borrower. The CCP samples individuals, not households or loans, so loans with coborrowers would be overrepresented by taking a simple average across observations. To correct for this issue, I weight each borrower by the fraction of their payments for which they are the sole borrower, such that the payments of a borrower who only has solo accounts gets weight 1 and a borrower who only has joint accounts gets weight $\frac{1}{2}$. That is, if a fraction j of a borrower's total mortgage payments are on an account for which they are a coborrower, the weight their payments get in the quarterly average is $1 - 0.5j$. This scheme provides equal weight to each borrowing unit, generally a household. I also drop observations with reported monthly mortgage payments in excess of $50,000 as potentially representing data errors. Using the weights to correct for the double-counting inherent to joint accounts, and adjusting for the CCP's coverage of 5 percent of all individuals with a social security number and a credit record, the data suggest an average of 42 million households with a mortgage in the U.S.

The time series of average monthly payments for households with a mortgage is shown in Figure 1, converted into real 2023 dollars. The average payment over the sample is $2,000, with considerable variance over the cycle. Some households hold multiple mortgages (for, e.g. investment or vacation properties) which pushes up the mean payment amount per borrower household relative to the mean payment per loan.

For mortgage rates, I use the 30-year, fixed-rate mortgage prime rate reported by Freddie Mac in their Primary Mortgage Market Survey (PMMS). The quarterly average of this series is plotted alongside the payments series in Figure 1. A positive relationship between average payments and mortgage rates, with a lag of a year or two, is visually apparent.

Mortgage originations come from data collected under the Home Mortgage Disclosure Act (HMDA). HMDA requires reporting on every mortgage application received for the majority of the market—estimates suggest HMDA's coverage is 90 percent of all originations (see Bhutta, Laufer and Ringo, 2017 or Johnson and Todd, 2019). Among other fields, lenders report loan amount, loan purpose, lien status, origination date, and number of units in the property. Between 1999 and 2022, HMDA lenders reported $712 billion in originated first liens on 1- to 4-unit dwellings per quarter, on average, in 2023 dollars.[10] Of this, $385

[10]Lien status and number of units were first reported in 2004, so figures for 1999-2003 include junior liens

billion was refinances and $327 billion was for home purchase or other purpose. I use the confidential supervisory version of HMDA, as origination date is not provided in the public version of the data.

For the ARM channel estimates, I use individual loan servicing data provided by ICE, McDash. ICE, McDash data cover the majority of residential mortgage servicing rights in the U.S. Weighted by scheduled payments, approximately 14 percent of mortgages had an adjustable rate during my sample period. ARMs were much more common in the early 2000's, however, and had declined to less than 10 percent of all scheduled mortgage payments by 2023. Dropping observations with scheduled payment greater than $50,000, I have approximately 910 million loan/month observations of ARMs from 2000 through 2023. Over the years of this sample, the average monthly payment for an ARM was $2063 in real 2023 dollars.

As I describe in Section 5.3, I estimate the effect of monetary policy shocks on monthly payments for individual ARMs. To keep these regressions tractable, I use a 1 percent random sample of ARMs in the ICE, McDash data. This sample consists of 205,974 loans and 9.1 million monthly observations over the years 2000 through 2023.

2.1 Monetary Policy Shocks

The current stance of monetary policy, as well as central banks' management of market expectations of its future path, are endogenous to the state of the economy and its outlook. The factors that determine aggregate mortgage debt service (e.g., the price of homes or households borrowing choices) are likewise endogenous, so a simple regression of payments on mortgage rates will be biased. To consistently estimate the effect of monetary policy decisions on mortgage payments, I require an exogenous source of variation in policy that is not driven by economic conditions.

Following Kuttner (2001) and Cochrane and Piazzesi (2002), a literature has developed isolating the surprise component of monetary policy decisions by measuring movements in

and multifamily originations. Prior to 2018, only loans for the purpose of home purchase, refinance or home improvement were reported under HMDA. In 2018 and after, loans for any purpose were reported if they were secured by a residence.

debt yields in a narrow time window around Federal Open Market Committee (FOMC) announcements. These movements represent the news shock associated with each FOMC announcement—the component of monetary policy decisions that was not already anticipated by markets, and which should be exogenous to prevailing market conditions.

FOMC announcements convey information about the future path of monetary policy as well as today's federal funds rate. The direction and magnitudes of the news shocks about current and future policy may not align—for example, the FOMC could decline to adjust the current rate while using surprisingly dovish language about its expectations and plans for the future. The pricing of longer-maturity debt depends more on these future expectations than on current short-term rates, so the reactions of interest rates may differ along the yield curve. Interest rates for 30-year, fixed rate mortgages (the typical home loan in the U.S.) are tied to longer-maturity yields. Forward guidance by the FOMC—and potentially LSAPs—are the method by which the Fed can affect mortgage payments, rather than short-term adjustments to the federal funds rate. For this reason I use the monetary policy shock series esitimated by Swanson (2021) as a measure of the exogenous component of FOMC decisions. Swanson (2021) uses differential movement in yields at different maturities around FOMC announcements to separately identify a time series of short-term federal funds monetary policy shocks from a series of forward guidance shocks that affect the long end of the yield curve. In addition, Swanson (2021) identifies a third component—an LSAP shock—based on additional variation in long-term yields that appeared with the advent of quantitative easing in 2009. These shock series cover 171 FOMC announcements from 1999 (when my CCP payments sample begins) through mid-2019.

A potential counter to the argument that the identified monetary policy shocks are exogenous is the possibility of a Fed information effect—FOMC announcements may reveal private information the Fed has about future economic developments and the market reaction reflects the digestion of this information as well as any updating of beliefs about the future path of policy. In the appendix, I show my results are robust to using the estimated monetary policy shock series of Lewis (2023), who partials out a Fed information effect by also considering the reaction of equity markets to FOMC announcements.

3 Monetary policy shocks and mortgage interest rates

As a first step to understanding how monetary policy affects mortgage debt service, I need to determine its effects on mortgage rates. To do this, I estimate the effects of the Swanson (2021) monetary policy shocks on the prime (mortgage) rate in the four years following each FOMC announcement using the local projections method of Jordà (2005). That is, for each shock S in week t, I estimate the effect on the prime rate r in week $t+k$ as:

$$r_{t+k} = \alpha_0^k + \alpha_1^k S_t + \sum_{l=1}^{L} \alpha_{1+l}^k r_{t-l} + \epsilon_t^k \tag{1}$$

sequentially for each $k = 1...208$. The shocks S are constructed to be exogenous, so α_1^k provides an estimate of the effect of a one-standard deviation contractionary shock on rates k weeks in the future. The exogeneity of S stems from their measurement as the information relevent to securities pricing unanticipated by markets. The anticipatable reaction by the FOMC to current or expected future economic developments should already be priced in, so news shocks contained in the announcements should reflect only idiosyncratic policy decisions exogenous to economic conditions. Ringo (2024) shows that the Swanson (2021) shocks are indeed uncorrelated with lagged values of mortgage rates, supporting the exogeneity argument.

Mortgage rates are strongly auto-correlated, so including lags of the outcome as controls (r_{t-l}) can improve efficiency. I choose the number of lags, L, to minimize Akaike's information criteria (AIC) for $k = 1$. In my sample, this was reached at $L = 2$.

Results from estimating equation 1 are presented in Figure 2. I begin with the forward guidance shocks, with results shown in the top panel. A one standard deviation contractionary shock raises the prime rate by an average of 7.4 basis points for three years following the announcement. In the third year, the effect diminishes before disappearing around three years post-shock. For comparison, Swanson (2021) reports that a one standard deviation forward guidance shock immediately raised the 5-year and 10-year treasury yields by about 5 basis points and 4 basis points, respectively. Note that in Figure 2, the effect on mortgage rates within the first few weeks post-shock is within this range, and subsequently increases over the ensuing years.

Results for federal funds and LSAP shocks are shown below in the middle and bottom panels of Figure 2. The federal funds shock never shows clear evidence of affecting mortgage rates. LSAP shocks do appear to raise rates notably for several weeks post announcement, but the effect quickly fades.[11] The transient effect of LSAP shocks on yields was also noted by Swanson (2021).

4 Estimating the effect of a monetary policy shock on payments

To see how much of these rate shocks passes through to mortgagor households, I next estimate the effect of a one standard deviation forward guidance shock on average payments. I again use the local projections estimator, repurposing equation 1 with the log of average monthly mortgage payments (p) as the outcome variable. The CCP data is reported quarterly, so I sum all forward guidance shocks within a quarter, q, to produce the net shock S for that quarter. I estimate:

$$p_{q+k} = \beta_0^k + \beta_1^k S_q + \sum_{l=1}^{L} \beta_{1+l}^k p_{q-l} + \nu_q^k \tag{2}$$

sequentially for $k = 1...16$. The AIC for this specification was minimized at $L = 5$, so I control for 5 lags of p.

Results are presented in the top panel of Figure 3. The effect of a tightening shock grows steadily through the first two years before levelling out in year 3 and then declining. At the three year mark, the point estimate suggests the average borrower is paying 0.89 percent more per month as a result of a one standard deviation positive forward guidance shock. As total annual mortgage payments were approximately \$1 trillion in real 2023 dollars across my sample, this implies a total increase in payments of \$9 billion per year, three years subsequent to the FOMC announcement. Extrapolating to a one percentage point shock,

[11] I re-sign the LSAP shocks so that a positive value corresponds to a monetary contraction for all three shock series.

these estimates suggest that the interest rate semi-elasticity of mortgage payments reaches about 10 if the adjustment to rates persists for three years.

A $9 billion increase in disposable income from a 7 basis point reduction in rates is quite large. As argued in the Introduction, if most or all of this marginal disposable income is consumed, the mortgage payment channel alone can explain a sizable fraction of the overall consumption response to monetary policy before considering indirect effects through, e.g., increased wages or wealth effects through higher house prices.

A semi-elasticity of 10 is also quite large given all the frictions to refinancing. Depending on the interest rates outstanding loans are paying, a wholesale reduction of rates for every existing borrower by one percentage point would reduce principal and interest payments by only a little more than 10 percent. As I will show in the next section, this surprising elasticity is mostly due to large volumes of inframarginal borrowing, as loans less than three years old have historically accounted for the substantial majority of outstanding U.S. mortgage debt (by dollars, although not by loan counts). Marginal refinances and ARM payments further enhance the elasticity of mortgage payments with respect to rates. Overall, the strength of the mortgage payment channel of monetary policy is almost as great in the U.S. as if changes in mortgage rates were fully passed through to all outstanding loans—albeit with a substantial lag of three years.

5 Evaluating the channels

The results from Section 4 indicate how much total mortgage payments change due to a forward guidance shock. However, this does not tell us the mechanism by which payments change. This section separately estimates the sizes of the inframarginal borrower, marginal borrower, ARM, and intensive margin borrowing channels.

5.1 A simple calculation of monetary policy pass-through via inframarginal borrowers

I begin with the inframarginal borrower channel. Using the data on mortgage origination flows from Section 2 and the estimates from Section 3, I can perform a simple back-of-the-envelope calculation to approximate how much a typical monetary policy shock would be expected to affect the mortgage payments of inframarginal borrowers.

Assume all borrowers get 30-year, fixed-rate mortgages. Let r be the prevailing prime rate which all borrowers pay. If a monetary policy decision decreases the prime rate to $r-\Delta$, then a loan of size M would see its monthly principal and interest payments decrease by:

$$M\left[\frac{\frac{r}{12}(1+\frac{r}{12})^{360}}{(1+\frac{r}{12})^{360}-1}-\frac{\frac{r-\Delta}{12}(1+\frac{r-\Delta}{12})^{360}}{(1+\frac{r-\Delta}{12})^{360}-1}\right] \tag{3}$$

relative to the counterfactual in which it was originated under the older, higher rate r.

The average prime rate over my sample period was 5.1 percent. Plugging this value in for r, and a Δ of 7.4 basis points (the effect of a one standard deviation forward guidance shock), inframarginal borrowers pay 0.0046 cents less per month for every dollar borrowed post-shock. With \$791 billion of new borrowing originated each quarter (adjusting for HMDA's incomplete coverage), new borrowers in the first quarter after the shock end up paying \$432 million less annually than they would have absent the shock.

In the next quarter, another \$791 billion is borrowed, and those borrowers are also subject to \$432 million less in annual payments. The first quarter's borrowers, meanwhile, continue to pay their lower, policy affected rates. After two quarters, therefore, the additional annual cost imposed on inframarginal borrowers is 2 × \$432 million, or \$864 million. This continues each quarter, until three years after the shock borrowers are collectively paying \$5.2 billion less per year.

Once the shock to interest rates, Δ, reverses, we should expect the effects on total payments to slowly fade as mortgages originated during the period of lower rates prepay.

A consideration of marginal borrowers naturally brings some asymmetry into the calculations above. For rate *cuts*, marginal borrowers are recruited from households that would not have borrowed under older, higher rates. The effects of monetary policy pass-through on inframarginal borrowers, as calculated in this section, and on marginal borrowers can therefore simply be added together. For rate *hikes*, however, marginal borrowers are those households that would have borrowed under the old rate but instead choose not to. Inframarginal borrowers are the remainder, so the estimated magnitude of the inframarginal borrower channel in response to a rate increase should be discounted by the fraction of borrowers that are marginal. For example, suppose households would borrow $\$M$ in new originations in the absence of any rate shock. If a rate hike caused x percent of these potential borrowers to instead take no action, new inframarginal borrowing would instead by $M\frac{100-x}{100}$. This value should be substituted into equation 3 in place of M.

The number of borrowers that are marginal to a given rate shock (x) depends on the size of the change in rates. As I show in Section 5.2, marginal borrowers make up only a few percent of overall borrowers for a typical forward guidance shock. The asymmetry they induce can therefore be ignored in the context of most monetary policy decisions. That is, x is small so $M\frac{100-x}{100} \approx M$ For very large rate hikes (i.e., those of multiple percentage points), a substantial fraction of borrowers would be marginal and the inframarginal channel should be attenuated accordingly. However, the relative rate-inelasticity of home purchase loan demand in particular limits this attenuation (see Sections 5.2 and 6). Home buyers' decisions to move are mostly attributable to factors such as getting a new job or having a baby, with interest rates playing only a small role.

5.2 Estimating the effect of a monetary policy shock on refinances and home purchase loans

To pin down the size of the marginal refinance channel, I again use a local projections estimator and the same Swanson (2021) forward guidance shock. The outcome variable this time is the log of the weekly number of refinance originations reported in HMDA. The AIC for this specification was minimized at $L = 12$, so I control for 12 lags of log refinance originations. Results are presented in the middle panel of Figure 3.

Refinance volumes respond to rate shocks just as expected, falling quickly as rates rise in response to the shock and staying depressed for two years. Then, as rates return to baseline in the third year (see Figure 2), refinances recover as well. Note that refinance volumes are plotted by week of origination. Given the typical lag of weeks or months between initial application (when the borrower's interest rate is usually locked in) and origination, it is expected that the full effect of the rate shock on originations would not be observed for a number of weeks following the FOMC announcement. This delay can explain the ramp-up of the effect observed over the first few weeks in the figure.

Over the first three years post-announcement, the average weekly effect of a one standard deviation forward guidance shock was to reduce refinance volumes by 3.5 percent. In my sample period, annual refinance originations reported in HMDA averaged 5.2 million per year. Again assuming HMDA has 90 percent market coverage, this implies 17.3 million refinance loans over a typical three year period, so a 3.5 percent decrease means a loss of 606,000 refinances. Estimates from the literature suggest that refinancing saves the average borrower around $3,000 per year (see Agarwal et al., 2023, Abel and Fuster, 2021, or Anenberg, Scharlemann and van Straelen, 2023, for example).[12] Together, these estimates imply a one standard deviation forward guidance shock increases mortgage payments by roughly $2 billion through the marginal refinance channel.

For the marginal purchase borrower channel, I repeat this estimator on a weekly series of the log number of home purchase originations reported in HMDA. The AIC for this specification was also minimized at $L = 12$, so I similarly control for 12 lags of log purchase originations. Results are presented in the top panel of Figure 4.

Unlike refinances, home purchase loans do not show much evidence of responding to monetary policy shocks. This is consistent with the literature that finds home buying demand

[12]Marginal refinancers experience much larger payment reductions from a rate cut than inframarginal borrowers because their counterfactual is the original rate on their loan, r^o, not the prevailing prime rate r. For a borrower induced into refinancing by a rate cut of size Δ, the total reduction in rate is therefore $r^o - (r - \Delta)$, whereas for inframargional borrowers the reduction in rate relative to the counterfactual of no shock is just Δ. The average refinancing over the last 25 years reduced the borrower's rate by roughly a full percentage point (Freddie Mac, 2021).

is rate-elastic only among a minority of the borrower population (i.e., lower-income, credit constrained households; see Bhutta and Ringo, 2021 and Ringo, 2024).

The role of marginal home purchase borrowers in the pass-through of monetary policy to mortgage payments is complicated, as some of these marginal borrowers will be first-time home buyers. These buyers are often leaving a rental unit, so the change in mortgage payments they bring does not translate straightforwardly into changes in household disposable income. However, the estimated number of marginal home purchase borrowers is small enough that these complications are negligible relative to the aggregate effect of monetary policy shocks on payments. See Appendix Section A.3 for a fuller discussion of the potential magnitude of sample selection bias from marginal first-time buyers.

5.3 Estimating the effect of a monetary policy shock on ARM payments

To determine the size of the ARM channel, I estimate the effect of a monetary policy shock on average ARM monthly payments. Unlike the CCP and HMDA, the ICE, McDash servicing data do not provide near-comprehensive coverage of the market. In particular, coverage was expanding rapidly in the early 2000s. A simple time series average of scheduled payments could therefore be contaminated by selection bias as more loans are brought into the sample.

Rather than estimating the effect of monetary policy shocks on an aggregate series, I therefore instead use a panel of loans and estimate the effect of shocks on payments within each individual loan. This should eliminate bias from larger or smaller loans entering the sample over time.

I again use a local projections estimator and the forward guidance shocks from Swanson (2021), this time run on a panel of individual loan data. Data are recorded at a monthly frequency, so I sum all forward guidance shocks within a month, m, to produce the net shock for that month. For each ARM i, the outcome variable is the log of the scheduled monthly payment, $p_{i,m}$. I estimate:

$$p_{i,m+k} = \delta_0^k + \delta_1^k S_m + \sum_{l=1}^{L} \delta_{1+l}^k p_{i,m-l} + \eta_{i,m}^k \tag{4}$$

Huber-White robust standard errors are clustered at the monthly level to account for correlated shocks across loans. The AIC for this specification was minimized at $L = 2$, so I control for 2 lags of log scheduled payments.

Results are presented in the bottom panel of Figure 3. The estimates are never statistically significant, but the point estimates suggest that ARM borrowers were paying 1.2 percent more on their loans three year after a one standard deviation forward guidance tightening shock. Note that the point estimates do not immediately lift off from zero after the FOMC announcement. ARMS are often structured for the rate to reset only once per year, and often begin with a multi-year fixed-rate period. Some delay in effect is therefore to be expected. Note also that ARM payments are tied to different portions of the yield curve than the prime mortgage rate is. Typical reference rates are 6-month or 1-year yields. The change in payments ARM borrowers experience in response to a given monetary policy shock would therefore be expected to be different than those of inframarginal fixed-rate mortgage borrowers.

The ICE, McDash data suggest ARMs represent 14 percent of all mortgage payments over the years 2000-2023. From the CCP, total mortgage payments averaged \$1 trillion per year in 2023 dollars. If we take the point estimates from the bottom panel of Figure 3 at face value (despite the wide confidence intervals), this implies a one standard deviation forward guidance shock increases total annual ARM payments by approximately \$2 billion.

A comparison of per-borrower effects for ARMs (1.2 percent) and the general mortgagor population (0.89 percent, as estimated in Section 4) at the three-year mark is also telling. Despite the widespread use of fixed-rate mortgages, the effect on the average borrower is 75 percent as large as on the average ARM borrower and appears with a similar lag. With the caveat that the ARM estimates in particular are imprecise, this suggests that the pass-through of monetary policy to mortgage payments is only modestly impaired by the idiosyncratic popularity of fixed-rate mortgages in the United States.

5.4 Estimating the effect of a monetary policy shock on loan size

The magnitude of the inframarginal channel calculated in Section 5.1 takes the volume of inframarginal borrowing as constant. However, borrowers that were inframarginal on the extensive margin (i.e., that were going to borrow regardless of the stance of monetary policy) may have been marginal on the intensive margin (i.e., their choice of how much to borrow depends on rates). To determine the size of the intensive margin channel, I run the local projections estimator using the weekly log of the average loan size of originations reported in HMDA as the outcome variable. The AIC for this specification was minimized at $L = 1$, so I control for 1 lag of log average loan size.

Results are presented in the bottom panel of Figure 4. The estimates through the first two years post-shock are very close to zero and far from statistically significant. The intensive margin channel apparently has a negligible effect on payments for years after a monetary policy shock. This finding is in line with a literature that finds very small rate elasticities of intensive margin borrowing (see DeFusco and Paciorek, 2017, Bhutta and Ringo, 2021 or Ringo, 2024 for example). Any intensive margin effect within borrowers that does exist may be additionally offset by a selection effect, as lower-income borrowers (who request smaller loan amounts) are more likely to drop out in response to a rate hike.

Interestingly, through the third and fourth years after the shock, the point estimates begin to trend more negative. While they are never statistically significant, by the fourth year they are large enough to have a substantial effect on total payments. This decline (which mostly occurs after the effect of the monetary policy shock on mortgage rates has dissipated) could be due to a delayed effect on house prices and hence purchase loan volumes. It could also help explain the declining effect of a monetary policy shock on total payments in the fourth year post-shock seen in the top panel of Figure 3.

Together, the effects through the ARM channel estimated in Section 5.3 ($2 billion), the effects through the marginal refinance channel estimated in Section 5.2 ($2 billion), and the effects through the inframarginal channel calculated in Section 5.1 ($5.2 billion) add up to close to the total effect on payments estimated in Section 4 ($9 billion) by year 3. This confluence was not imposed by the estimators, since none of the estimates provided here in Section 5 use the estimates or time series of average payments from Section 4. Consideration of the marginal purchase borrower channel and intensive margin channel does little to change this summation, since their estimated magnitudes are small-to-zero.

6 State dependence

The effects of a monetary policy shock estimated and calculated so far in this paper represent average treatment effects across the years of the estimation sample. Within this time period, conditions have varied considerably, however. The total effect via the inframarginal borrower channel should scale directly with origination volumes (see equation 3), so effectiveness depends on future mortgage demand. This section documents the extent to which that demand, and hence the strength of the inframarginal borrower channel, has varied over time.

In Figure 5, I show for each quarter since 1999Q1 the total dollar volume of mortgage originations over the next three years reported in HMDA in real 2023 dollars. Mortgages are split into refinance and other purpose (mostly home purchase). The three-year peaks for total lending (for all purposes) in the refinance booms of 2003 and the Covid era were around $13 trillion, while future borrowing stayed below $8 billion from 2006 through 2018. Refinance lending in particular is volatile. The ratio of refinance volumes between the most- and least-active three-year periods was 3.1, while the ratio for the most- and least-active three-year periods for all borrowing was 2.4. Even in times when refinancing dries up, the less-volatile demand for home purchase borrowing provides a floor to total borrowing. For example, in the aftermath of the taper tantrum in 2013 the next three years of mortgage borrowing never fell below $6 billion dollars. This ensures some continued effectiveness of the inframarginal borrower channel even when the refinancing channel shuts down.

The continued effectiveness of the mortgage payment channel even when borrowers are generally out of the money is hinted at by the Covid-era experience. As can been seen in Figure 1, mortgage rates fell rapidly from 2018 through 2020. With a delay of somewhat over a year, mortgage payments fell as well. In 2022 and 2023, mortgage rates reversed trend and began rising even more rapidly than they had fallen as the Federal Reserve committed to series of rate hikes in response to persistent inflation. Again with some delay, mortgage payments began swiftly rising, despite the refinance market being totally moribund. The extent to which payments will eventually rise due to this tightening cycle is unknown as of the time of this writing, with the prime rate only hitting its peak in late 2023. However, it appears that higher rates are still being passed through to mortgage debt service in meaningful magnitudes.

7 Conclusion

I show that monetary policy has a powerful effect on household disposable income net of debt service obligations through affecting mortgage payments. However, most of this effect does not occur through the marginal refinancing channel that has received much attention in the literature. Instead, inframarginal borrowers are the most important single channel. Borrowers whose mortgage has an adjustable rate have also been an important channel in the past, although their prevalence is reduced in today's market. The large volume of originations (in the average year, Americans originate new mortgage volumes equivalent to one quarter of the total outstanding volume of mortgage) means that the distinction between fixed-rate and adjustable-rate mortgage regimes may not be as stark as generally perceived. Extrapolating the estimated effect sizes in this paper, mortgage rate swings that are well within historical experience can generate changes in payments worth hundreds of billions of dollars per year.

8 Bibliography

Abel, Joshua, and Andreas Fuster. 2021. "How do mortgage refinances affect debt, default, and spending? Evidence from HARP." *American Economic Journal: Macroeconomics*, 13(2): 254–291.

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tim Landvoigt, Tomasz Piskorski, Amit Seru, and Vincent Yao. 2023. "Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinance program." *The Review of Economic Studies*, 90(2): 499–537.

Amromin, Gene, Neil Bhutta, and Benjamin J Keys. 2020. "Refinancing, monetary policy, and the credit cycle." *Annual Review of Financial Economics*, 12: 67–93.

Anenberg, Elliot, Therese C Scharlemann, and Eileen van Straelen. 2023. "Borrowing and Spending in the Money: Debt Substitution and the Cash-out Refinance Channel of Monetary Policy."

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra. 2019. "Regional heterogeneity and the refinancing channel of monetary policy." *The Quarterly Journal of Economics*, 134(1): 109–183.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra. 2021. "Mortgage prepayment and path-dependent effects of monetary policy." *American Economic Review*, 111(9): 2829–2878.

Bernanke, Ben. 2007. "The financial accelerator and the credit channel." Board of Governors of the Federal Reserve System (US) Speech.

Best, Michael Carlos, James S Cloyne, Ethan Ilzetzki, and Henrik J Kleven. 2020. "Estimating the elasticity of intertemporal substitution using mortgage notches." *The Review of Economic Studies*, 87(2): 656–690.

Bhutta, Neil, and Benjamin J Keys. 2016. "Interest rates and equity extraction during the housing boom." *American Economic Review*, 106(7): 1742–1774.

Bhutta, Neil, and Daniel Ringo. 2021. "The effect of interest rates on home buying: Evidence from a shock to mortgage insurance premiums." *Journal of Monetary Economics*, 118: 195–211.

Bhutta, Neil, Steven Laufer, and Daniel R Ringo. 2017. "Residential mortgage lending in 2016: Evidence from the Home Mortgage Disclosure Act data." *Federal Reserve Bulletin*, 103: 1.

Calza, Alessandro, Tommaso Monacelli, and Livio Stracca. 2013. "Housing finance and monetary policy." *Journal of the European Economic Association*, 11(suppl_1): 101–122.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary policy when households have debt: new evidence on the transmission mechanism." *The Review of Economic Studies*, 87(1): 102–129.

Cochrane, John H, and Monika Piazzesi. 2002. "The fed and interest rates-a high-frequency identification." *American economic review*, 92(2): 90–95.

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

DeFusco, Anthony A, and Andrew Paciorek. 2017. "The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit." *American Economic Journal: Economic Policy*, 9(1): 210–40.

Di Maggio, Marco, Amir Kermani, and Christopher J Palmer. 2020. "How quantitative easing works: Evidence on the refinancing channel." *The Review of Economic Studies*, 87(3): 1498–1528.

Di Maggio, Marco, Amir Kermani, Benjamin J Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao. 2017. "Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging." *American Economic Review*, 107(11): 3550–88.

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong. 2022. "State-dependent effects of monetary policy: The refinancing channel." *American Economic Review*, 112(3): 721–761.

Federal Reserve Board. 2023. "Financial Accounts of the United States - Z.1." *public release.*

Flodén, Martin, Matilda Kilström, Jósef Sigurdsson, and Roine Vestman. 2021. "Household debt and monetary policy: Revealing the cash-flow channel." *The Economic Journal*, 131(636): 1742–1771.

Freddie Mac. 2021. "Refinance Trends in the First Half of 2021." Research Note.

Gorea, Denis, Oleksiy Kryvtsov, and Marianna Kudlyak. 2022. "House Price Responses to Monetary Policy Surprises: Evidence from the US Listings Data."

Hamilton, James D. 2008. "Daily monetary policy shocks and new home sales." *Journal of Monetary Economics*, 55(7): 1171–1190.

Hofmann, Boris, and Gert Peersman. 2017. "Is there a debt service channel of monetary transmission?" *BIS Quarterly Review, December.*

Jappelli, Tullio, and Luigi Pistaferri. 2010. "The consumption response to income changes." *Annu. Rev. Econ.*, 2(1): 479–506.

Johnson, Kevin, and Richard M. Todd. 2019. "The Value of HMDA Coverage of Home Lending in Rural Areas and Indian Country."

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American economic review*, 95(1): 161–182.

Kim, Youngju, and Hyunjoon Lim. 2020. "Transmission of monetary policy in times of high household debt." *Journal of Macroeconomics*, 63: 103168.

Kuttner, Kenneth N. 2001. "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of monetary economics*, 47(3): 523–544.

Lewis, Daniel J. 2023. "Announcement-Specific Decompositions of Unconventional Monetary Policy Shocks and Their Effects." *The Review of Economics and Statistics*, 1–46.

Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2021. "Local projection inference is simpler and more robust than you think." *Econometrica*, 89(4): 1789–1823.

Ringo, Daniel. 2024. "Monetary policy and home buying inequality." *Review of Economics and Statistics*, 1–46.

Rubio, Margarita. 2011. "Fixed- and Variable-Rate Mortgages, Business Cycles, and Monetary Policy." *Journal of Money, Credit and Banking*, 43(4): 657–688.

Swanson, Eric T. 2021. "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets." *Journal of Monetary Economics*, 118: 32–53.

Wong, Arlene. 2019. "Refinancing and the transmission of monetary policy to consumption." *Unpublished manuscript*, 20.

Figure 1: Average Monthly Mortgage Payment and Prime Mortgage Rate, by Quarter

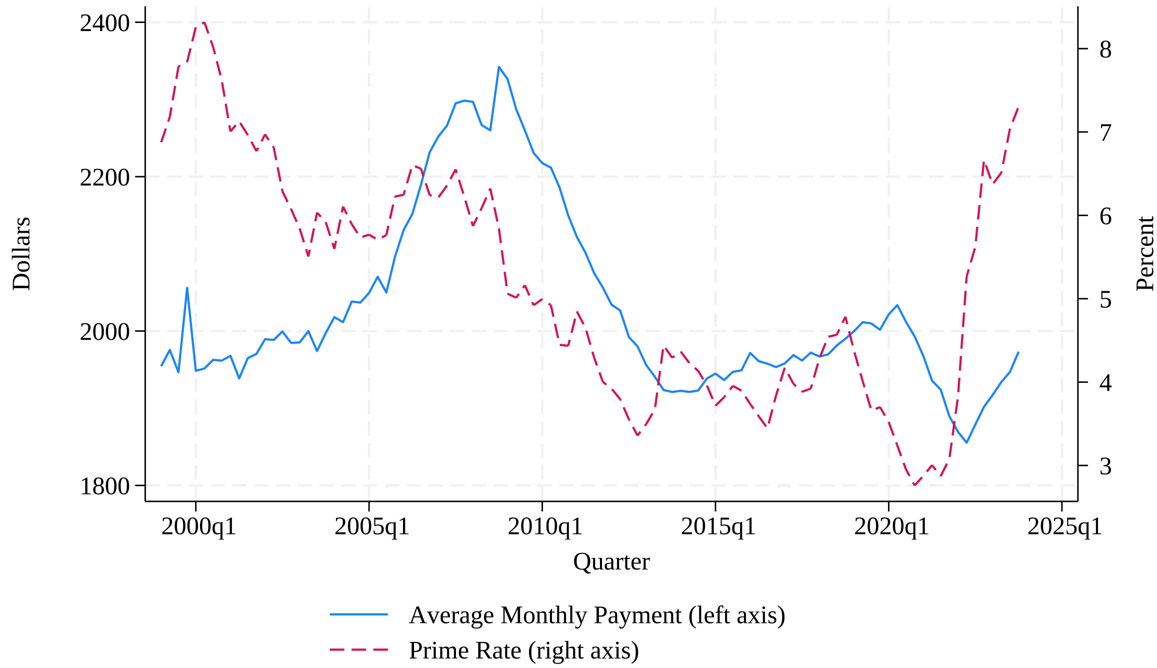

Note: The solid line shows the average monthly payment on first-lien mortgages for households with a mortgage, by quarter, in real 2023 dollars. Based on data from the Federal Reserve Bank of New York/Equifax Consumer Credit Panel. The dashed line shows the quarterly average prime rate for 30 year, fixed-rate mortgages reported in Freddie Mac's Primary Mortgage Market Survey.

Figure 2: Effect of Monetary Policy Shocks on Mortgage Rates, by Week

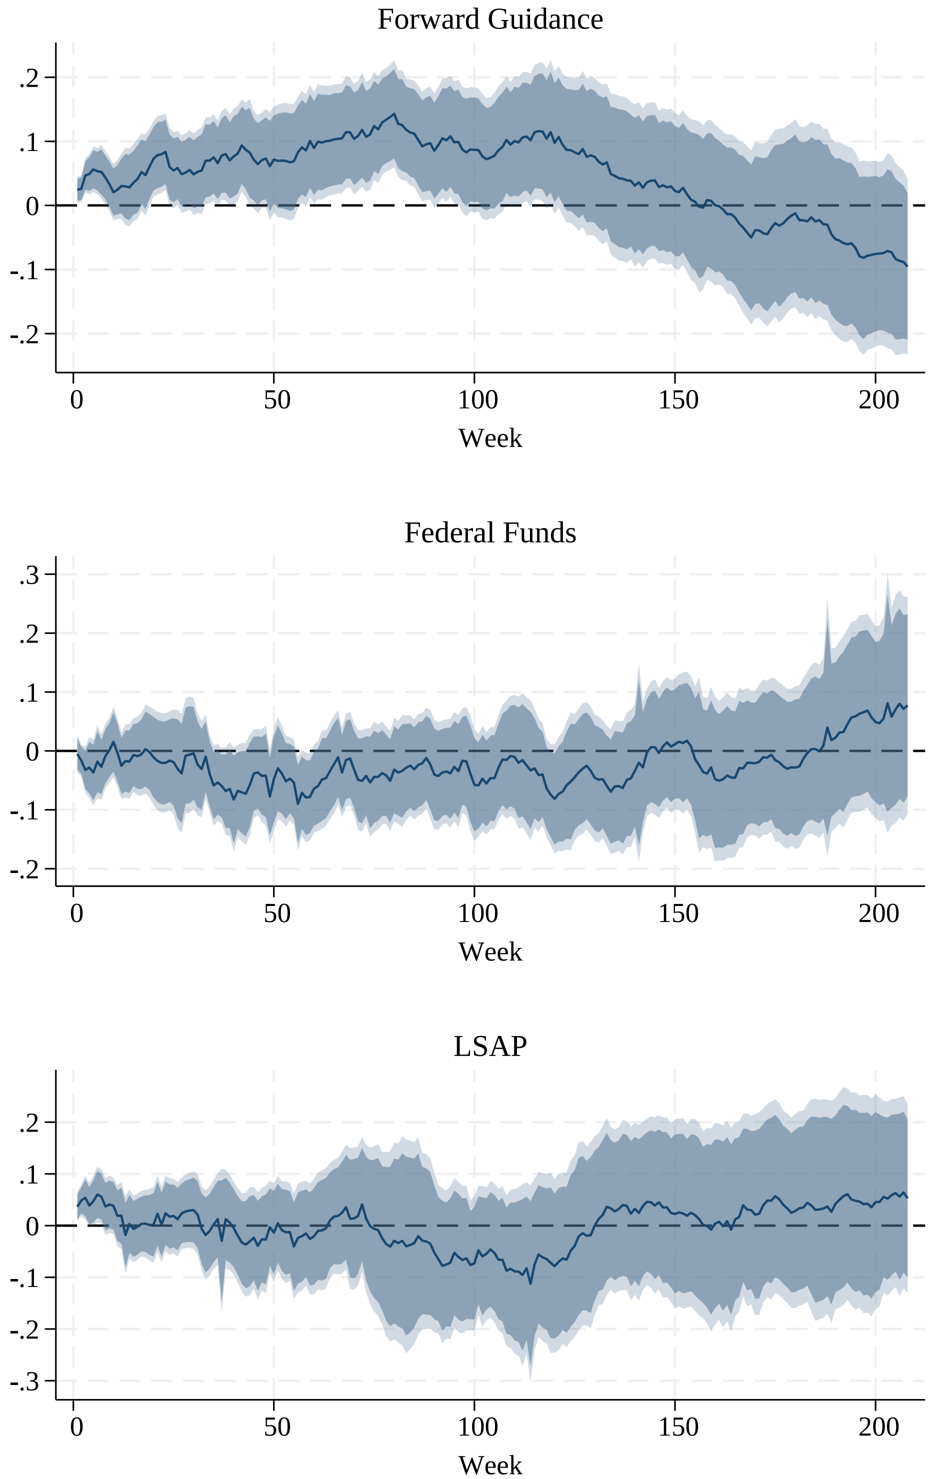

Note: Estimated effect of a one standard deviation contractionary monetary policy shock in week $t = 0$ on the weekly prime mortgage rate (measured in percentage points) for the 4 years after the shock. Values of the monetary policy shocks, split into LSAP, forward guidance, and federal funds shocks, are sourced from Swanson (2021). Estimates are obtained using the local projection method of Jordà (2005). Shaded areas represent the 95 and 90 percent confidence intervals, calculated from Huber-White heteroskedasticity robust standard errors, following Montiel Olea and Plagborg-Møller (2021).

Figure 3: Effect of Monetary Policy Shocks on Mortgage Payments and Refinances

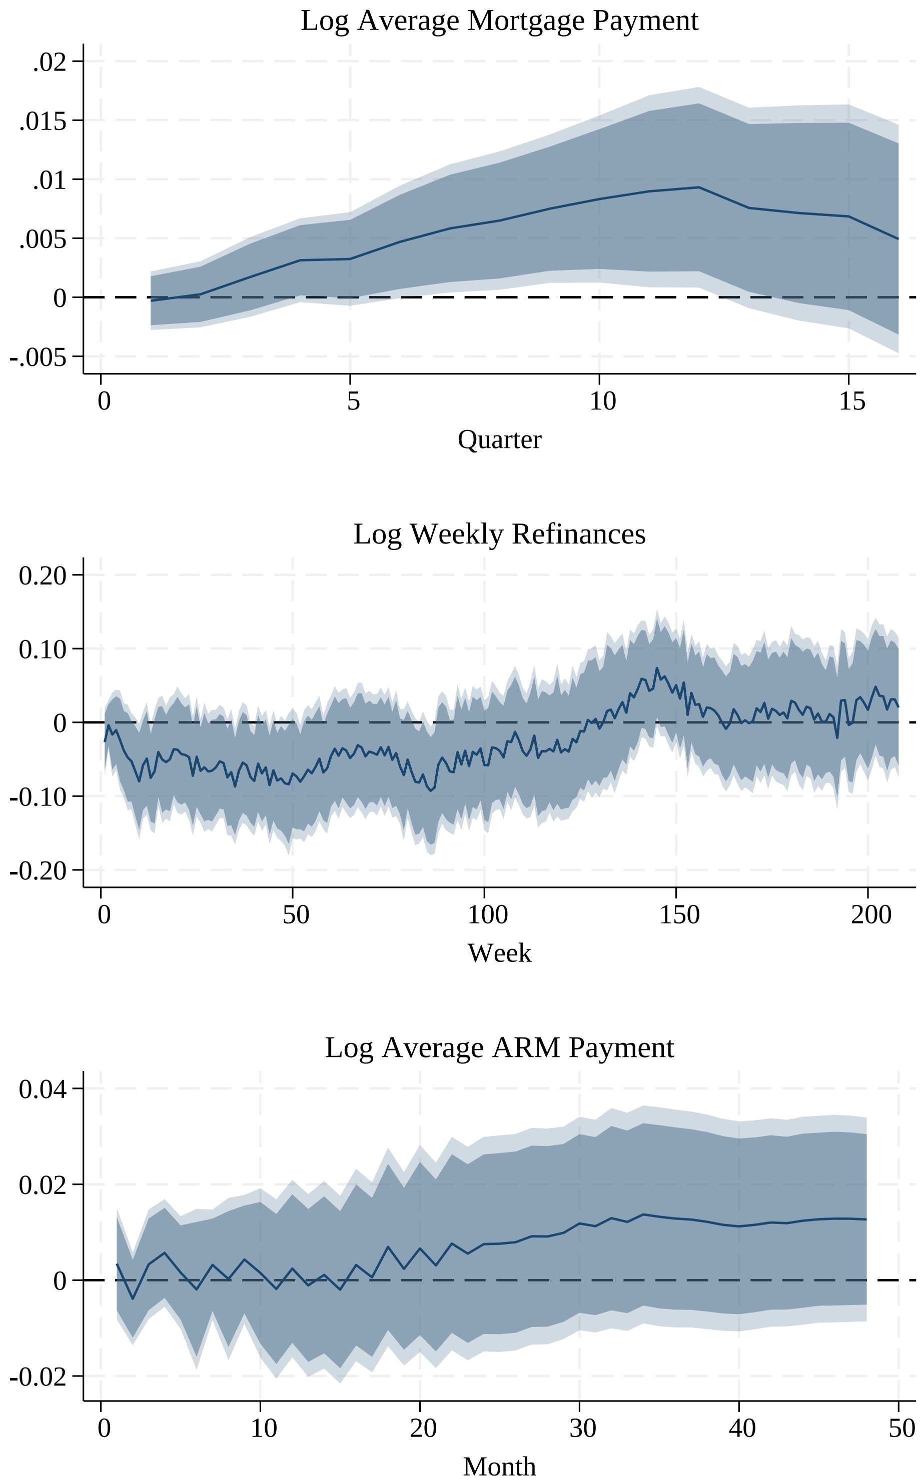

Note: Estimated effect of a one standard deviation contractionary forward guidance monetary policy shock in period $t = 0$. Top panel shows effects on log average monthly payments of all outstanding mortgages, by quarter. Middle panel shows effects on log weekly refinance originations. Bottom panel shows effects on the log average monthly payments of ARMs. Values of the monetary policy shocks are sourced from Swanson (2021). Estimates are obtained using the local projection method of Jordà (2005). Shaded areas represent the 95 and 90 percent confidence intervals, calculated from Huber-White heteroskedasticity robust standard errors, following Montiel Olea and Plagborg-Møller (2021). For the bottom panel, standard errors are clustered at the monthly level.

Figure 4: Effect of Monetary Policy Shocks on Purchase Originations and Loan Amounts

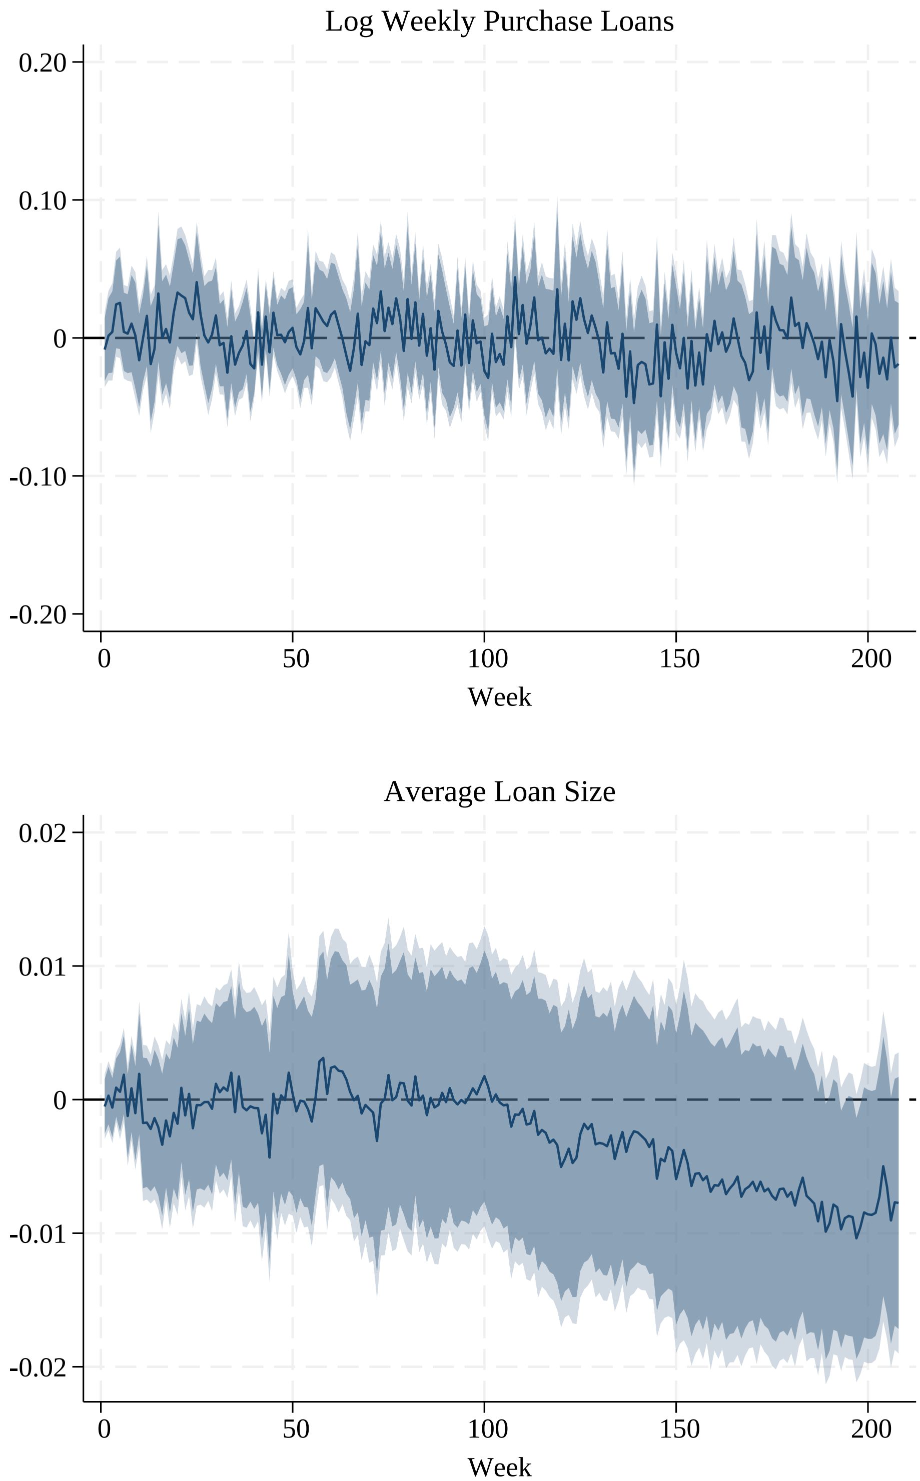

Note: Estimated effect of a one standard deviation contractionary forward guidance monetary policy shock in period $t = 0$. Top panel shows effects on log weekly home purchase originations. Bottom panel shows effects on log weekly average loan size of originations. Values of the monetary policy shocks are sourced from Swanson (2021). Estimates are obtained using the local projection method of Jordà (2005). Shaded areas represent the 95 and 90 percent confidence intervals, calculated from Huber-White heteroskedasticity robust standard errors, following Montiel Olea and Plagborg-Møller (2021).

Figure 5: Total Originations Over the Next Three Years

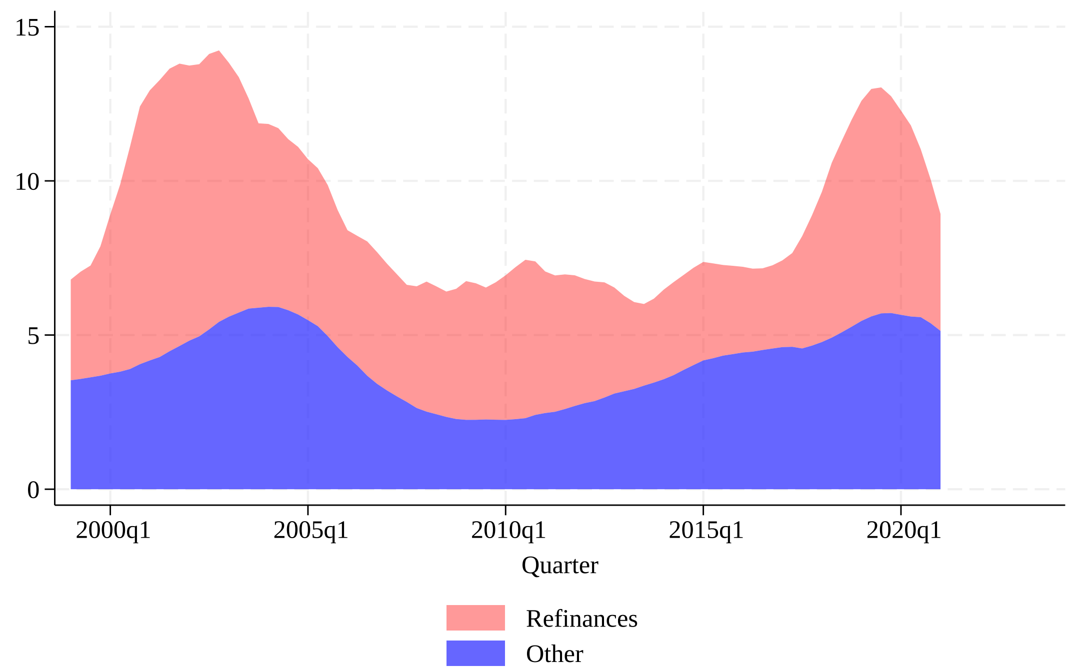

Note: The total dollar volume of originations reported in the Home Mortgage Disclosure Act data over the subsequent three years, by quarter. Volumes, reported in trillions of real 2023 dollars, are cumulative across the "Refinance" and "Other" categories, such that the level of the total shaded area indicates the total amount borrowed regardless of purpose.

Internet Appendix for

"Inframarginal Borrowers and the Mortgage Payment Channel of Monetary Policy"

Daniel Ringo

July 30, 2024

A Appendix

A.1 Alternative estimates of monetary policy shocks

In this section I show my empirical results are robust to using alterative monetary policy shock series estimated by Lewis (2023).

Like Swanson (2021), Lewis (2023) estimates distinct components of the surprise information contained in FOMC announcements, broken out into components that affect the short-term federal funds rate and longer term rates via forward guidance and large-scale asset purchases (LSAP). He also extracts a Fed information effect component which should not properly be considered a monetary policy shock. While the end goal is similar to that of Swanson (2021), Lewis (2023) uses a different identification strategy to tease out these components, tracking the evolution of rates over the rest of the afternoon post-announcement to see how markets absorb the information non-instantaneously. This different method leads to different measures of the shocks: the correlation between the forward guidance shocks estimated by Swanson (2021) and Lewis (2023) is only 0.61. This difference in measured shocks could potentially lead to different estimates of their effects.

I redo the analysis from Sections 3, 4, and 5 using the Lewis (2023) shocks instead. For ease of comparison with the baseline results, I normalize the shocks Lewis (2023) reports to have a unit standard deviation.[1] I begin with the effects of the various shocks on mortgage rates, shown in Figure A.1. Similarly to the Swanson (2021) shocks shown in Figure 2, I find that the forward guidance shock of Lewis (2023) has a lasting effect on mortgage rates. A one standard deviation forward guidance shock is estimated to raise rates by 6.3 basis points on average over the three years post-announcement. The federal funds shocks and LSAP shocks do not show a lasting effect on mortgage rates.

With a slightly smaller effect on rates, we would expect a proportionately smaller effect on payments. Plugging a Δ of 6.3 basis points in equation 3, I calculate that a one standard deviation Lewis (2023) forward guidance shock should increase payments by inframarginal

[1]Following Swanson (2021), I scale the federal funds shocks to their standard deviation over the pre-2009 period, and the LSAP shocks to their standard deviation over the post-2009 period to reflect the respective environments in which standard policy and unconventional policy (the zero lower bound era) were operating.

borrowers by $4.4 billion per year at the three-year mark post-announcement.

Next, I re-estimate the effects on total payment, marginal borrowers, ARM payments, and loan size shown in Figures 3 and 4. Results are presented in Figures A.2 and A.3. The effect on total payments, shown in the top panel, is very similar to that found using the Swanson (2021) shocks, although slightly diminished (as expected) to match the smaller effect on rates. Three years after the FOMC announcement, a one standard deviation tightening is estimated to increase average payments by 0.76 percent, or $7.6 billion per year.

The estimated effect of a one standard deviation Lewis (2023) forward guidance shock on the number of refinances is shown in the middle panel of Figure A.2. Despite the smaller estimated effect on rates, the estimated effect on refinance volumes is no lower than that estimated using the Swanson (2021) shocks. A 6.3 basis point rate increase is now estimated to reduce refinancing volumes by an average of 3.9 percent over the three years post-announcement. At $3000 of annual savings per refinance, this similarly delivers about $2 billion in additional payments through the marginal refinance channel. I show the estimated effect on ARM payments in the bottom panel of Figure A.2. The point estimates do not suggest any effect on ARM payments (although, as with the Swanson (2021) shocks, standard errors are large enough that I cannot reject a meaningful size change).

Estimated effects on home purchase borrowing and average loan size are shown in Figure A.3. As with the Swanson (2021) shocks, effects on home purchase borrowing are negligible. However, the intensive margin estimates are non-significantly *positive*.

A.2 Comparison to Cloyne, Ferreira and Surico (2020)

In this section I compare my estimates of the effect of a monetary policy shock on average mortgage payments to those of Cloyne, Ferreira and Surico (2020).

There are some differences between Cloyne, Ferreira and Surico (2020) and this paper in the time period studied, the monetary policy shock series used, and the specific estimator. However, the key difference is their use of the Consumer Expenditure Survey (CEX) to construct a time series of average mortgage payments. The CEX interviews a few thousand mortgagor households per quarter and asks them questions including the amount paid on fixed-rate mortgages last month. Interviewees rotate out of the panel after 4 quarters. The CCP data I use samples approximately 3 million mortgage borrowers (borrowers randomly selected based on social security number) who remain in the panel until death, and derives the payment information from the scheduled amounts loan servicers report to the credit bureau.

The time series properties of mortgage payments in the CCP and CEX are noticeably different. I construct a CEX-based analog to my CCP mortgage payment time series from 1999-2022. Following Cloyne, Ferreira and Surico (2020), I include only households that held a mortgage in all four consecutive quarters in which they were interviewed. This leaves an average of approximately 2500 mortgagor households sampled per quarter. The average payments constructed in this manner (weighted for representativeness using the CEX's provided weights) are shown in Figure A.4, in real 2023 dollars, with the CCP payments shown alongside for comparison. As can be seen, the CEX series shows considerably more volatility quarter-to-quarter, and somewhat different cyclical trends as well. The one-quarter autocorrelation coefficient for the CCP payments data is 0.98, while that for the CEX data is only 0.9, suggesting the presence of statistical noise in the CEX data.

When I re-estimate equation 2 on the CEX data (again using the Swanson 2021 forward guidance shocks), I no longer find a statistically significant relationship. Results are presented in Figure A.5. While the point estimates are different from Cloyne, Ferreira and Surico (2020) (unsurprising given the different time period, shock series, and specific estimator), the general lack of evidence of a major positive effect of rates on payments is consistent with their study.

The most likely explanation for the difference in results obtained using the CCP and the CEX data is that measurement error in the CEX obscures the actual relationship between monetary policy shocks and mortgage payments. The CEX payments data exhibit considerable quarter-to-quarter volatility that likely reflects noise from small sample sizes, changing cohorts of interviewees, and interviewee response error. The identifying variation from estimated monetary policy shocks is only a few basis points of rate changes every FOMC announcement, so a small amount of noise could easily overwhelm the signal from the actual effects of these shocks. Consider also the results presented in Section 5 of the present paper. These were estimated without using the CCP payment time series, and yet sum up to a similar overall effect of monetary policy shocks on payments. If the CEX-based results are to be believed, there would have to be a countervailing force leading to approximately $9 billion in payment *decreases* caused by a 7 basis point rate increase to make the net payment change essentially zero in light of the evidence from Section 5. In the absence of a theory for such a large countervailing force, the CCP-based results seem more reliable.

A.3 Marginal first-time buyers and sample selection bias

In Section 4, I estimated the effect of a monetary policy shock on the average payments of mortgage borrowers. This outcome could be affected both by mortgagor households owing different payments due to the shock (the mortgage payment channel) and by a change in the composition of mortgagors—i.e., the entrance (or abstention) of marginal first-time buyers. Previous literature has found that the rate elasticity of home buying is much greater among lower-income households (Bhutta and Ringo, 2021, Ringo, 2024), suggesting that marginal first time buyers have lower payments than the average mortgagor. This means that sample selection bias will tend to exaggerate my estimates of the magnitude of the mortgage payment channel.

In this section I show that marginal first-time home buyers are too few in number to materially bias my estimates of the mortgage payment channel of monetary policy. Section 5.2 suggests the number of marginal home buyers is small, but does not provide a precise estimate. Turning to the literature, Bhutta and Ringo (2021) and Ringo (2024) both estimate a rate semi-elasticity of home buying of about 20 for the population of below-average income home buyers. To be conservative, I assume that above-average income buyers have an elasticity of zero to maximize the potential sample selection bias. An overall rate semi-elasticity of 10 is well within the 95 percent confidence interval estimated in Section 5.2.

First-time home buyers represent 35 percent of home purchase borrowers by count, and are similarly rate-elastic as repeat home buyers (Ringo, 2024). With 11.7 million purchase originations over an average 3-year period, 4.1 million can be expected to be first time buyers. If 2.05 million are below-average income, a rate semi-elasticity of 20 implies the 7.4 basis point shock estimated in Section 3 would stop approximately 30,000 lower-income households from becoming first-time home buyers over a three-year period. These households would have represented 0.066 percent of the 46 million mortgagor households in the U.S. Even if the monthly payments of a marginal first-time buyer were half of those of an average buyer, their presence would bias the estimate of the mortgage payment channel up by only 3.3 basis points (or $400 million), about 4 percent of the total estimated effect.

Figure A.1: Effect of Lewis (2023) Monetary Policy Shocks on Mortgage Rates, by Week

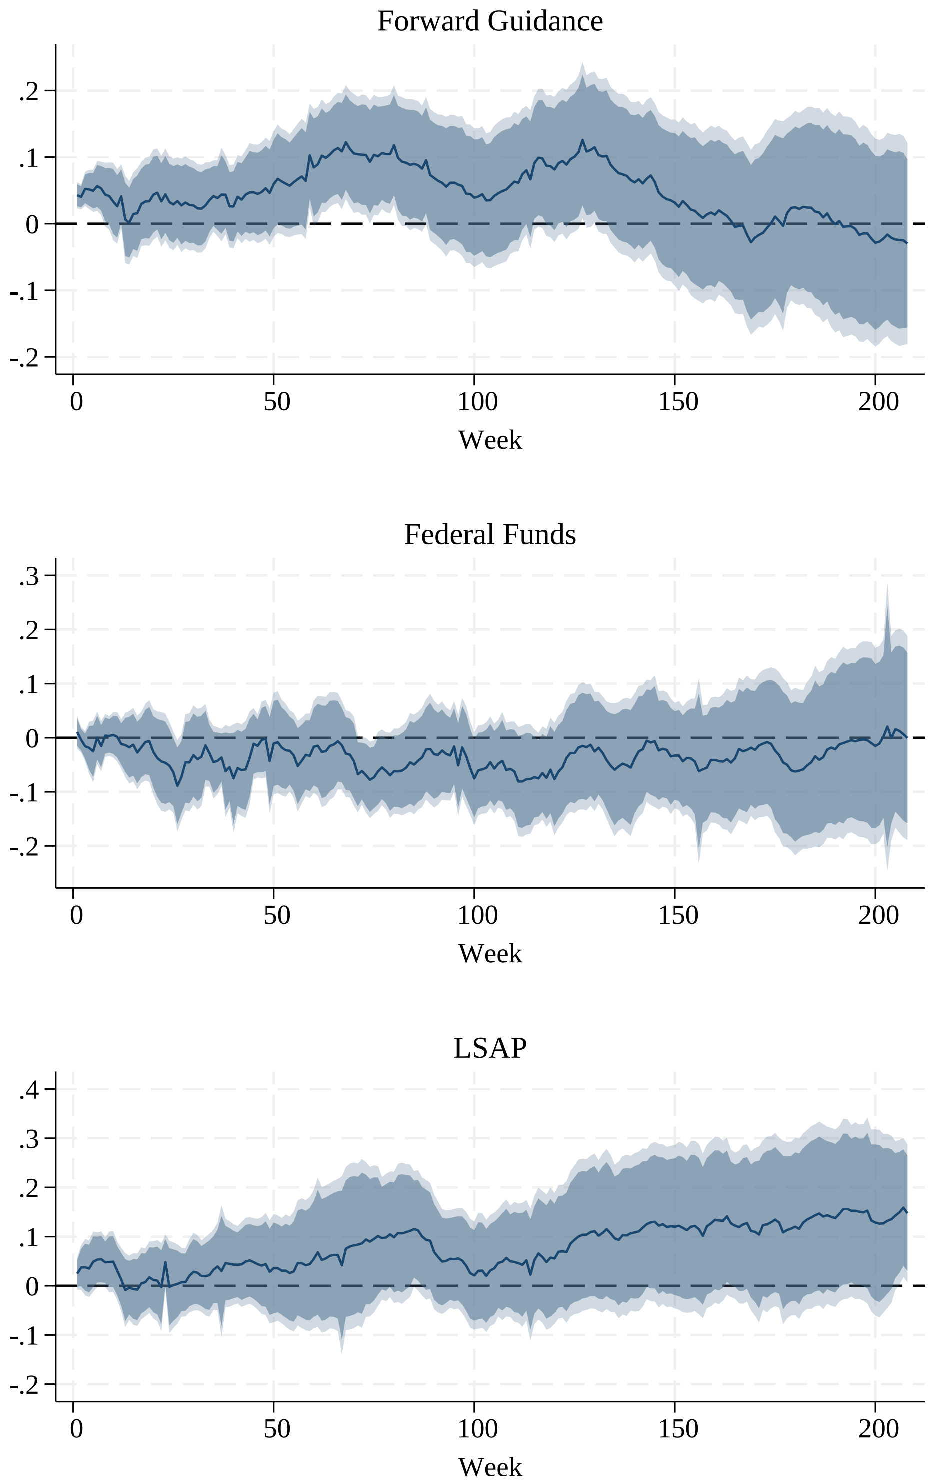

Note: Estimated effect of a one standard deviation contractionary monetary policy shock in week $t = 0$ on the weekly prime mortgage rate (measured in percentage points) for the 4 years after the shock. Values of the monetary policy shocks, split into LSAP, forward guidance, and federal funds shocks, are sourced from Lewis (2023). Estimates are obtained using the local projection method of Jordà (2005). Shaded areas represent the 95 and 90 percent confidence intervals, calculated from Huber-White heteroskedasticity robust standard errors, following Montiel Olea and Plagborg-Møller (2021).

Figure A.2: Effect of Lewis (2023) Monetary Policy Shocks on Mortgage Payments and Refinances

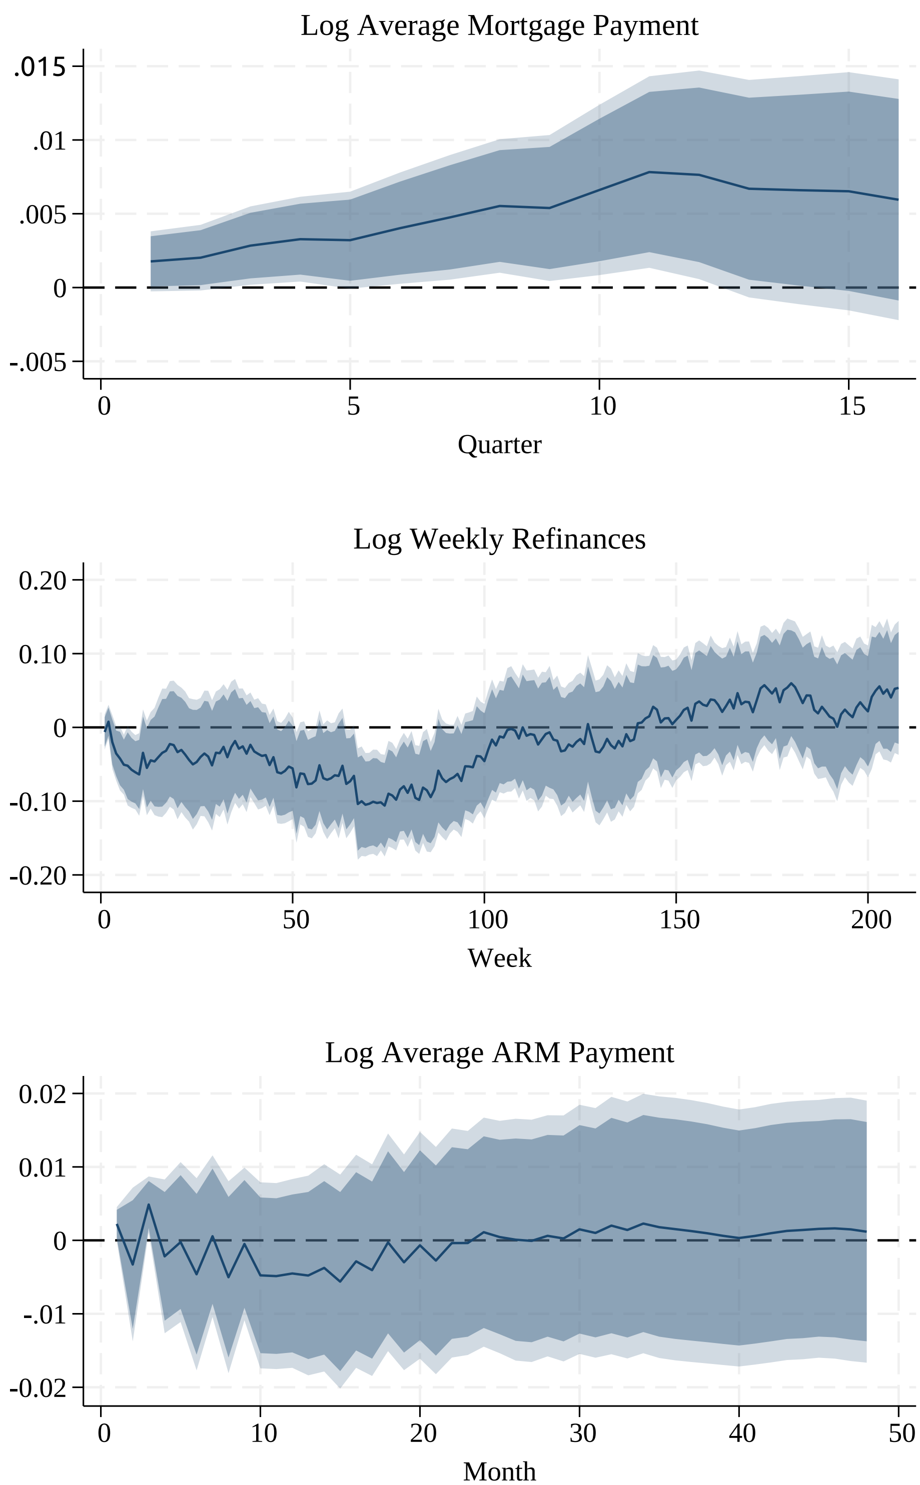

Note: Estimated effect of a one standard deviation contractionary forward guidance monetary policy shock in period $t = 0$. Top panel shows effects on log average monthly payments of all outstanding mortgages, by quarter. Middle panel shows effects on log weekly refinance originations. Bottom panel shows effects on the log average monthly payments of ARMs. Values of the monetary policy shocks are sourced from Lewis (2023). Estimates are obtained using the local projection method of Jordà (2005). Shaded areas represent the 95 and 90 percent confidence intervals, calculated from Huber-White heteroskedasticity robust standard errors, following Montiel Olea and Plagborg-Møller (2021). For the bottom panel, standard errors are clustered at the monthly level.

Figure A.3: Effect of Lewis (2023) Monetary Policy Shocks on Purchase Originations and Loan Amounts

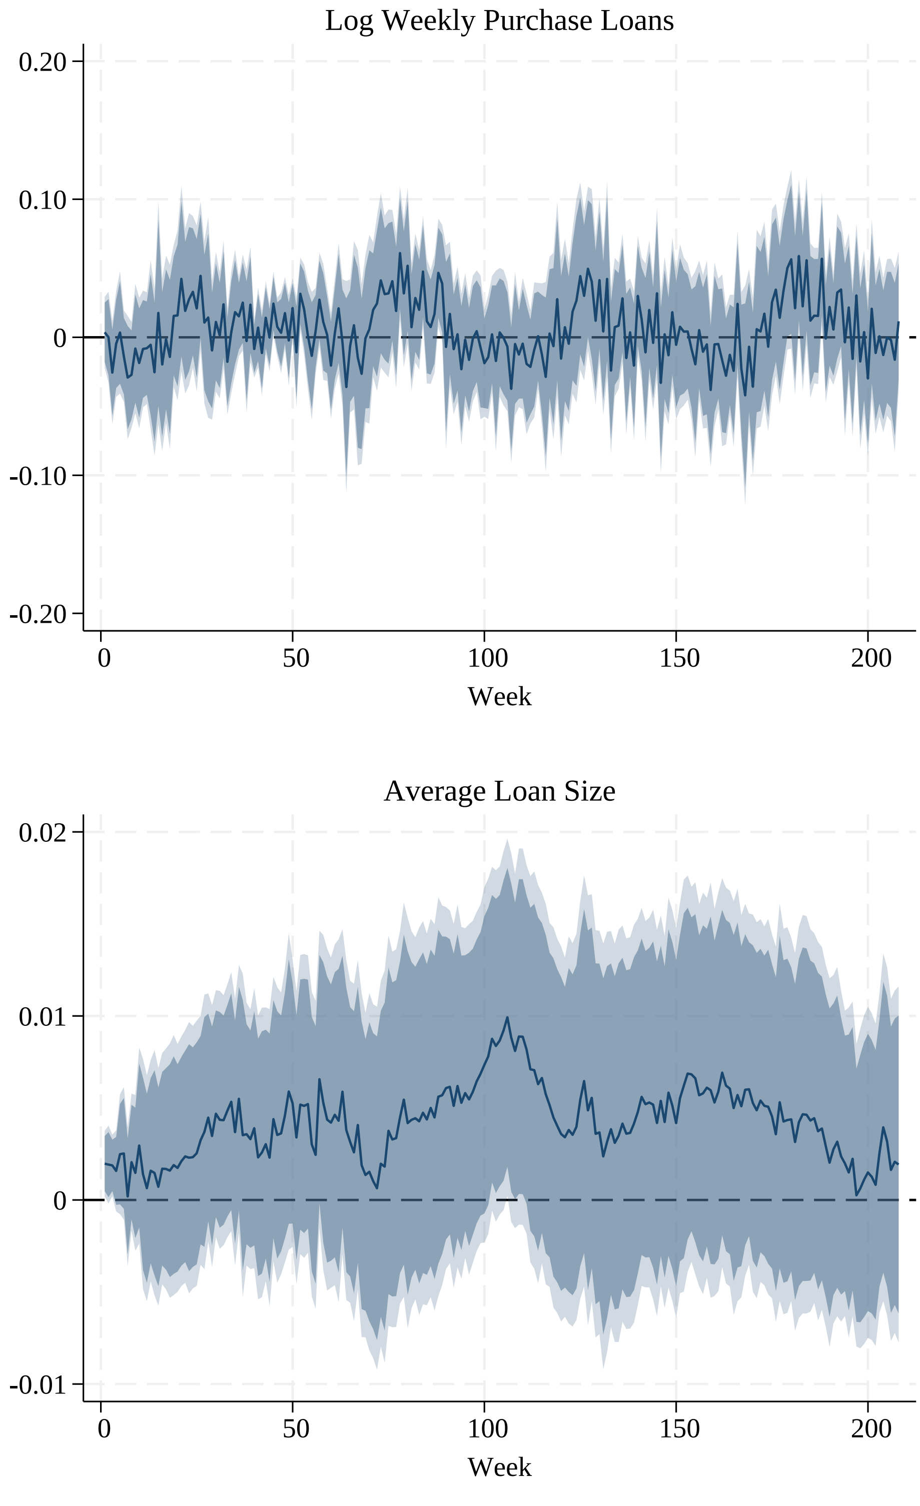

Note: Estimated effect of a one standard deviation contractionary forward guidance monetary policy shock in period $t = 0$. Top panel shows effects on log weekly home purchase originations. Bottom panel shows effects on log weekly average loan size of originations. Values of the monetary policy shocks are sourced from Swanson (2021). Estimates are obtained using the local projection method of Jordà (2005). Shaded areas represent the 95 and 90 percent confidence intervals, calculated from Huber-White heteroskedasticity robust standard errors, following Montiel Olea and Plagborg-Møller (2021).

Figure A.4: Average Monthly Mortgage Payment in the CCP and CEX, by Quarter

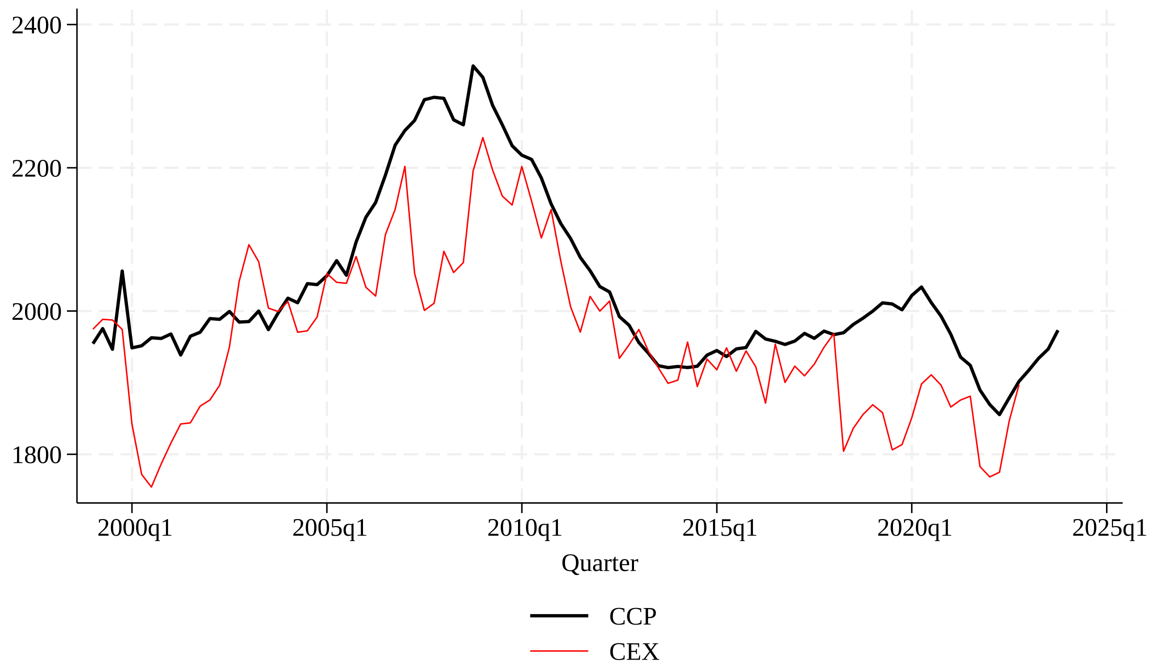

Note: The thicker black line shows the average monthly payment on first-lien mortgages for households with a mortgage, by quarter, in real 2023 dollars, as calculated from the Federal Reserve Bank of New York/Equifax Consumer Credit Panel (CCP). The thinner red line shows the average monthly payment on fixed-rate mortgages for households with such a mortgage, by quarter, in real 2023 dollars, as calculated from the Consumer Expenditure Survey (CEX).

Figure A.5: Effect of Monetary Policy Shocks on Mortgage Payments in CEX

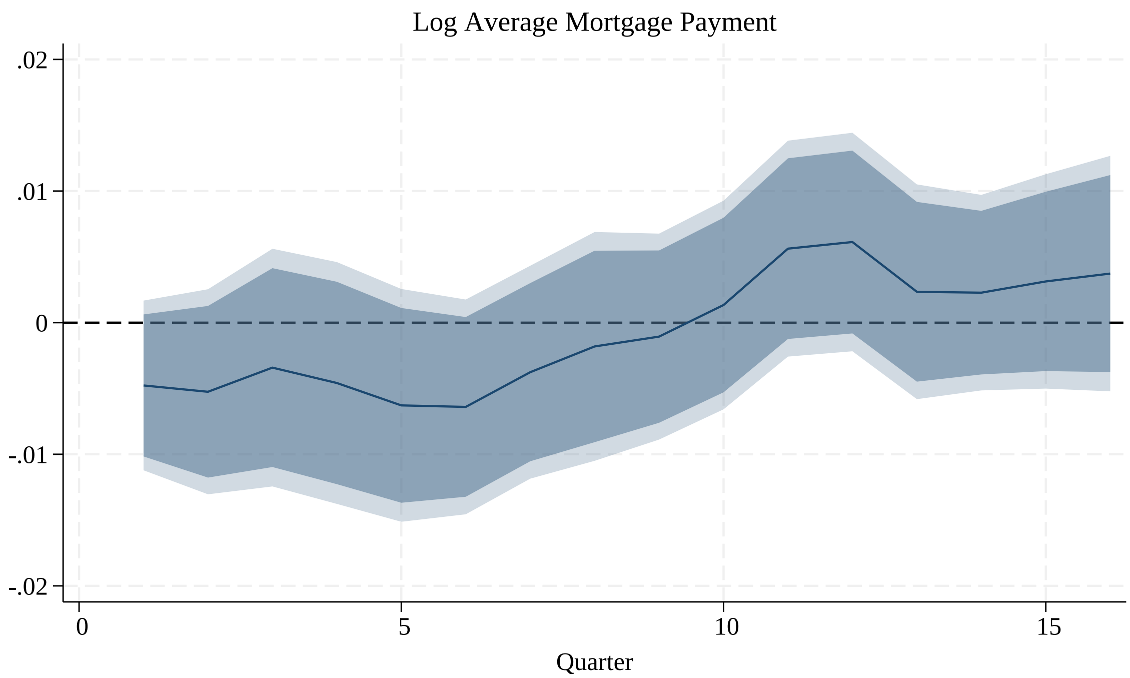

Note: Estimated effect of a one standard deviation contractionary forward guidance monetary policy shock in period $t = 0$ on monthly fixed-rate mortgage payments, as calculated from the Consumer Expenditure Survey (CEX), 1999-2022. Shaded areas represent the 95 and 90 percent confidence intervals