Are Millennials Ruining Monetary Policy?
Student Loan Debt and Monetary Policy Transmission

Molly Shatto*

(Click here for the most recent version)

November 7, 2024

Abstract

Has the significant increase in fixed-rate, non-refinanceable student loan debt in the United States over the last 20 years affected the transmission of monetary policy to household consumption? Several studies have emphasized the importance of refinancing as a key mechanism through which households fund consumption in response to expansionary monetary policy. In this paper, I estimate consumption responses to expansionary monetary policy shocks for households with different types of debt to study the implications of the increased share of student loan debt in the household balance sheet. I find that households with fixed-rate, non-refinanceable student loan debt have a muted consumption response compared to mortgagors and households with variable-rate student loan debt. Mortgagors and variable-rate student loan holders have significantly larger consumption responses, approximately $4000, or 40% more over four years. I then calculate the implied effect on average household consumption responses given the composition of the household balance sheet in 2007 versus 2019. I find that the reallocation of household debt has reduced average monetary policy responses by around 10% two years following an expansionary monetary policy shock.

*Department of Economics, UC San Diego, email: mshatto@ucsd.edu.

I would like to thank Valerie Ramey, Johannes Wieland, Juan Herreño, Nir Jaimovich, and Munseob Lee for their invaluable guidance and support. I am also thankful to Ali Uppal, Erica Chuang, and participants at the Second Women in Central Banking Workshop and the UCSD Macro Seminar for useful comments and suggestions.

1 Introduction

Over the last 20 years in the United States, unsecured, fixed-rate student loan debt has more than doubled, rising from 4% of total household debt in 2007 to over 10% by 2020. During this period balances have nearly quadrupled, with borrowers carrying an average of $40,000 in student loan debt in 2020. The rise in student loan debt has shifted the composition of the household balance sheet away from refinanceable mortgage debt and toward non-refinanceable[1] student loan debt. The secular rise of student loan debt may interfere with the refinancing mechanism, an essential channel through which households fund consumption in response to expansionary monetary policy (Wong 2021; Beraja et al. 2019; Berger et al. 2021; Eichenbaum et al. 2022).

This paper shows that the increase in student loan debt over the last 20 years has attenuated the transmission of monetary policy to household consumption. I use detailed quarterly consumer credit panel data from the California Policy Lab (UC-CCP) to estimate differential monetary policy responses for households with student loans and mortgages. I find that households with fixed-rate student loan debt have smaller consumption responses compared to those with mortgages. I also exploit a change in U.S. student loan policy that moved from variable interest rates to fixed interest rates in the Federal Direct Student Loan program. I show that borrowers with variable-rate student loans respond more to monetary policy shocks than those with fixed-rate student loans. These results highlight the importance of the refinancing channel, and suggests that the differing responses between mortgagors and fixed-rate student loan borrowers are also likely driven by refinancing and not due to differences between these groups. I then provide a back-of-the-envelope calculation to illustrate that these effects are quantitatively important. Using my estimated household responses, I find that the changing composition of household debt has resulted in a decrease in the average household consumption response of approximately 10%. This may significantly impact monetary policy effectiveness going forward if trends continue.

Rising tuition costs in the United States, paired with an increasing reliance on student loans, have led to a significant increase in student loan debt across demographic groups. The share of households with student loan debt increased from 12% to over 22% between 2007 and 2019, while the average student loan balance for borrowers has grown from $21,000 to $40,000 over the same period. Student loans are now the second largest debt category, surpassing motor vehicle and credit card debt. At the same time, there has been a fall in the share of households with mortgage debt from 49% at the peak in 2007 to just over 40% in 2019.

Determining the implications of these trends for monetary policy transmission is challenging. A

[1] To refinance federal student loans, individuals must go through private lenders. This is not generally a viable option for most households, as private student loans are subject to risk-based pricing, so rates can be significantly higher, and individuals lose most of the benefits of the federal student loan program, such as relatively generous income repayment plans, forbearance, and forgiveness for select professions.

key constraint is finding data that tracks individual consumption and contains reliable information on household balance sheet items. The benefit of using credit bureau data is the detailed information on household debt for individuals over time, along with detailed loan information, which can be used to construct measures of spending in crucial consumption categories, such as vehicle purchases. The measures of household debt from the UC-CCP corresponds closely with aggregate trends in both mortgage and student loan debt. I can infer the interest rate type for specific student loans using the history of interest rates in the student loan program, a feature that is not available in other higher frequency datasets with both balance sheet information and consumption, such as the Consumer Expenditure Survey. This added benefit allows me to test the importance of the refinancing mechanism more directly.

While consumption is not observed directly in the UC-CCP data, I construct a proxy of household-level consumption using auto loan originations following Aruoba et al. (2022). I find that this auto loan proxy variable from credit bureau data closely matches aggregate spending on vehicles from the Bureau of Economic Analysis (BEA). Therefore, changes in these loans can approximate increased spending on cars and other vehicles, which account for a sizable portion of consumer durable spending, and while only a subset of total household spending, consumer durables are particularly responsive to monetary policy (McKay and Wieland 2021).

To measure the effect of student loan and mortgage debt on household responses to monetary policy, I estimate individual-level local projections (LP) (Jordà 2005), in which the future values of consumption are regressed on current changes in interest rates interacted with an indicator for debt status. This allows me to estimate impulse response functions (IRFs) for individuals with different types of debt. I isolate the heterogeneous consumption responses for households with different types of debt by taking the coefficients on the interaction terms from these regressions to capture the dynamic effect of monetary policy on household consumption at each horizon.

Monetary policy is conducted in response to current and future expected conditions of the economy, which poses issues for casual identification. To identify the causal effects of interest rate changes induced by monetary policy, I use monetary policy shocks to isolate exogenous movements in the real interest rate. I take the monetary policy shock series from Bauer and Swanson (2023), who estimate monetary policy surprises by taking high-frequency interest rate changes around the Federal Reserve's Federal Open Market Committee (FOMC) announcements. These shocks work well over the relatively short time frame of my panel (2004-2019). The identifying assumption for my results is that there are no additional variables that are driving differential responses to the monetary policy shocks that are correlated with debt status, such as demographic differences between individuals with different types of debt. I therefore include a limited set of controls, including age and age interacted with the monetary policy shocks to account for differential monetary policy responses over the life-cycle.

The main specification tracks the additional consumption response of households with fixed-rate student loans and mortgage debt over the non-debtor[2] response. I find that the response to a one-standard-deviation expansionary monetary policy shock for individuals with fixed-rate student debt is only slightly lower than the non-debtor response, and this effect dissipates entirely after four years, resulting in a cumulative response that is not significantly different from the non-debtor response of $10,000 after 16 quarters. On the other hand, mortgagors will, on average, spend an additional $4000 on vehicles over four years, 40% more than the non-debtors.

To more directly test the affect of fixed-rate versus variable-rate debt, I exploit a 2006 policy change in the federal student loan program to compare households with variable-rate and fixed-rate student loan debt directly. I find that households with variable-rate student loans also have a significantly higher response, spending an additional $4000 over four years, almost identical to the mortgage response. Both of these results suggest that the refinancing mechanism is an essential driver of monetary policy responses.

Next, I present further evidence in support of the refinancing and interest rate exposure mechanisms and address any threats to identification. First, I assess the importance of the level of debt in driving the differential responses that I observe. The level of debt should correspond to the strength of the refinancing channel as individuals with large balances should see larger decreases in interest rate payments in response to lower rates, and thus increased cash flow. To test this assumption, I estimate additional specifications that directly test the importance of the level of debt. I find that the size of debt is a crucial driver of the household response for those with mortgages and variable-rate student loans, as the consumption response is increasing in the level for both. Conversely, the level of fixed-rate student loan debt does not appear to affect the consumption response. This is consistent with the refinancing mechanism, as higher levels of debt will increase the incentives of refinancing. I also show that mortgage refinancing increases and mortgage payments decrease in response to expansionary monetary policy shocks in the credit panel data, consistent with the existing literature.

Second, I test the robustness of my results to the inclusion of age controls.[3] As debt is not randomly assigned in the population, an ideal empirical design would include a number of individual and household-level fixed effects to account for selection into different loans. One limitation of the data is the overall lack of demographic variables. To assess how significant this concern is for my results, I estimate the main specification with and without the available age variables and find that the results are not sensitive to the inclusion of age controls. I also test the use of auto loan originations as the primary measure of consumption by looking at alternate consumption measures and find that the results do not significantly change.

[2] Here, non-debtor is defined as households without mortgage or student loan debt. They may hold other types of debt.

[3] I show, using data from the Survey of Consumer Finances, that age is the most significant characteristic where households with mortgages and student loans differ.

Finally, I construct a back-of-the-envelope calculation to quantify the implied effect of the changing balance sheet on average monetary policy responses and use a "Three-Agent" New Keynesian Model (ThANK)[4] to show that this back-of-the-envelope calculation can be informative for the aggregate monetary policy response. First, I weight the estimated impulse response functions by the share of households with each type of debt in 2007 and 2019, respectively, to generate average consumption responses and compare the weighted IRFs. These estimates suggest a decrease in the average household response of approximately 10% after two years. I then provide further context for the back-of-the-envelope calculation, using a stylized two-period ThANK model to examine how the micro-level estimates of the household consumption response aggregate to macro-level monetary policy responses. I find that under standard assumptions, the back-of-the-envelope is a good approximation of the aggregate effect. This suggests that the rise in student loan debt has significantly impacted the transmission of monetary policy to household consumption.

My empirical work relates to a growing theoretical literature that calls for the inclusion of heterogeneity in macroeconomic models, laid out extensively in the heterogeneous agent (HANK and TANK) literature (Kaplan et al. 2018; Ahn et al. 2018; Luetticke 2021; Debortoli and Galí 2017; Galí et al. 2007; Bilbiie 2008). Specifically, my paper builds on the work of Auclert (2019), who shows that the ability to adjust loan terms in response to changes in interest rates matters for the transmission of monetary policy, which Auclert terms the interest-rate exposure channel. I directly test the importance of fixed versus variable-rate debt in the transmission of monetary policy in the U.S. context, providing an empirical test of this important theoretical transmission mechanism and contributing to the growing literature on how monetary policy is transmitted through household debt (Slacalek et al. 2020; Jappelli and Scognamiglio 2018; Flodén et al. 2021).

This literature highlights the importance of this net interest rate exposure channel. In the U.S. context, where most mortgages are fixed rate, the literature has focused on refinancing as the primary driver of the consumption response in the U.S. (Wong 2021; Berger et al. 2021; Beraja et al. 2019; Eichenbaum et al. 2022). Cloyne et al. (2020) also find that mortgagors drive most of the aggregate monetary policy response in the U.S. and U.K., although the authors focus on the differences across the housing tenure status in MPCs and not on refinancing. My research adds to this literature by comparing the mortgage response to the response of those with student loan debt to measure the importance of this interest rate exposure channel for the transmission of monetary policy and to understand the implications of the increasing trend in student loan debt. I verify the the importance of mortgage refinancing, using credit panel data, which, to my knowledge, has not previously been used in this setting. I then provide evidence that those with fixed-rate student loan debt do not respond as strongly

[4] I extend the standard two-period version by including two hand-to-mouth agents, each of which holds exogenous debt with either variable-rate interest payments or fixed-rate interest payments, along with the Ricardian agent.

to expansionary monetary, suggesting that the rise in student loan debt has broader implications for effective monetary policy.

My work also relates to several studies in the household finance literature examining the connection between rising student loan debt and the housing market, and how the trends in student loan debt and mortgage debt may be inherently linked (Mezza et al. 2020; Amromin et al. 2017; Black et al. 2023; Cooper and Wang 2014; Houle and Berger 2015). This is still an active area of research with many questions still left unsettled, which may be important mechanisms for my setting.

In the context of monetary policy, mortgage debt and refinancing play an essential role in the transmission of monetary policy to real economic activity. The rise in student loan debt and the changing composition of household balance sheets may limit the ability of households to take advantage of favorable interest rates through refinancing, thus muting the impact of monetary policy.

The rest of the paper will proceed as follows: In section 2, I discuss the composition of household balance sheets and the relevant empirical setting. In section 3, I describe the data. I then lay out the estimation strategy and results in section 4 and present evidence of refinancing in section 5. In section 6, I then present a simple model to help contextualize my results and then conclude.

2 Empirical Setting

In this section, I lay out the relevant details for my empirical setting. I show how the composition of the household balance sheet has changed over time and the underlying dynamics of household debt that explain the relevant trends. I then briefly discuss the history of student loan policy with a particular focus on how interest rates have evolved over time and how they relate to other similar loan types.

2.1 Household Balance Sheet

For decades, mortgage debt has been the most important component of household debt. Recently, student loan debt has become a more prominent component, particularly for young households, as the share of households with student loans and average balances have increased relative to other types of debt. Figure 1 plots the trend in both student loan debt as a share of total household debt (left panel) and mortgage debt as a share of total household debt (right panel). Mortgage debt has fallen as a share of total household debt from its peak of 78% in 2007 to about 67% of total household debt in 2020. This decrease in the share of mortgage debt has been matched by a sharp increase in the share of student loan debt, which increased from around 4% in 2007 to over 10% of total household debt in 2020.

While Figure 1 shows that mortgage debt and student loan debt have been on offsetting trends, mortgage debt remains the single largest debt category, making up just over two-thirds of total household

Figure 1: Student Loan and Mortgage Debt as a Share of Total Household Debt, 2007-2019

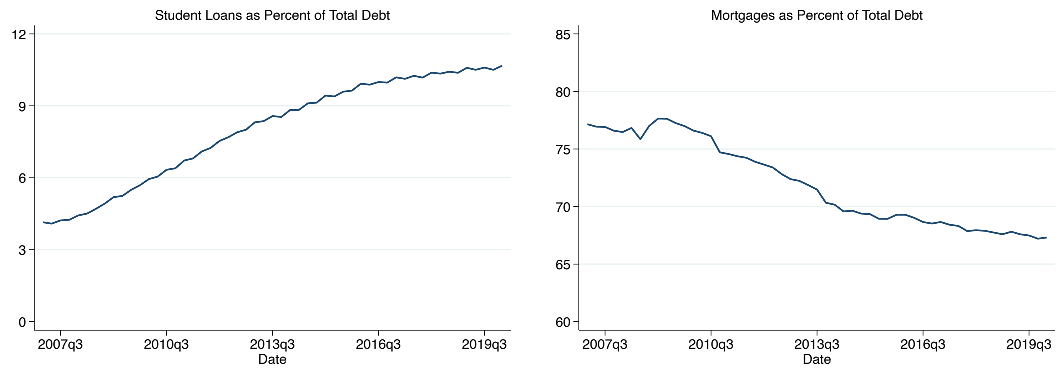

Notes: The figures show the quarterly change in the percent of total household debt that is student loan debt (left panel) and mortgage debt (right panel). The vertical axis is percent and the horizontal axis is quarters between 2007q1 and 2019q4. Source: Federal Reserve Board; Z.1 Financial Accounts of the United States, and G.19 Consumer Credit Release.

debt in 2020.[5] However, student loans over this 13-year period have passed auto loans and credit cards to become the second largest source of debt in the United States, topping 1.76 trillion dollars by the beginning of 2023.[6] Other loans, including credit cards and auto loans, stayed relatively constant throughout this period.[7]

For a clearer picture of the underlying dynamics behind these trends in household debt, in Appendix A, I include figures showing the share of households with student loan and mortgage debt, and average and median balances for borrowers of each type of debt. The share of households with student debt has more than doubled between 1989 and 2019 and the average balance has nearly quadrupled. On the other hand, the share of households with mortgage debt has fallen from its peak during the housing bubble in the mid-2000s. While the average student loan balance has grown consistently, the average mortgage balance has less than doubled over this period and has remained fairly constant since the Great Recession.

2.2 Student Loan Details

One way in which debt is important for the transmission of monetary policy is that changes in interest rates can change payments on debt. If interest payments decrease in response to monetary policy shocks, this opens up additional funds that can be used to finance consumption. Fixed-rate debt or non-refinanceable debt will shut down this interest rate channel, as interest payments on fixed-rate loans will not respond to interest rate changes.

Student loan debt, relative to other types of debt, is unique in that most of these loans are provided

[5] In Appendix A, Figure 10 plots the composition of total household debt over the 13 years from 2007 to 2020.
[6] Source: Federal Reserve Board, G19 Consumer Credit Outstanding
[7] See Appendix A for individual plots of changing shares.

through the federal government, and interest rates on all loans are set by Congress annually.[8][9] The early federal student loan program began in 1965 and provided subsidized loans through partner lenders, backed by the federal government. In 1992, with the re-authorization of the 1965 Higher Education Act, the direct student loan program was piloted. This new direct loan program switched, for the first time, to variable-rate student loans.[10] Some years later, interest rates in the direct loan program were reverted back to fixed rate, taking effect in 2006 (Delisle 2012). While the rates do vary from year to year, an individual's rate is fixed once the loan is taken out. Figure 2 shows the history of interest rates in the direct lending program since 2006.

Figure 2: Federal Student Loan Interest Rate History SY2006/07 - SY2024/25

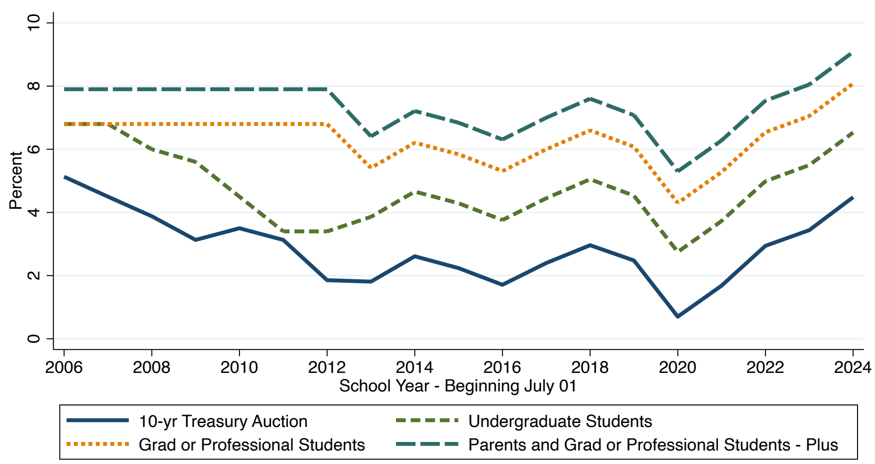

Notes: The figure shows the history of interest rates in the Federal Student Loan program, along with the 10-year Treasury Auction rate. The horizontal axis is school years, and the vertical axis is percent. Source: the U.S. Department of Education and the U.S. Department of the Treasury.

These rates are generally favorable, as rates are determined by loan type and not credit score, therefore students do not face risked-based market pricing as they would for private student loans, or for other kinds of unsecured loans. Table 1 shows these rates for the 2023-24 and 2024-25 school years, along with average mortgage rates, and quotes for some private student loans and personal loans.

There is almost no ability to refinance or discharge these loans. It is not possible to refinance federal loans with government lenders. You can consolidate your loans at the average of your current interest rates but this will not lower rates. To refinance student loans, individuals must go through private

[8]Details on the history of student loan interest rates and servicers is available through the U.S. Department of Education.

[9]These rates are determined using the 10-year Treasury note auction yield plus an add-on margin, and they apply to new federal loans disbursed from July 1st of the current year to July 1st of the following year (Cox 2017). The add-on margin is calculated so the government breaks even.

[10]The variable rates reset once a year based on the interest rates on short-term U.S. Treasury securities plus 3.1 percentage points, capped at 9.0 percent (Delisle (2012)).

Table 1: Interest Rates on Different Types of Loans

Type of Loan	Interest Rate
Undergraduate	5.50%
Graduate & Professional	7.05%
Parent, Graduate & Professional PLUS	8.05%

(a) Panel A: Federal Student Loan Rates SY23-24.

Type of Loan	Interest Rate
30-Yr FRM	6.89%
15-Yr FRM	6.17%
5/1-Yr ARM	6.64%

(b) Panel B: Average Mortgage Rates as of July 2024.

Type of Loan	Interest Rate
Private Fixed Rate Student Loan	4.19% - 14.83%
Private Variable Rate Student Loan	5.74% - 15.86%
Personal Loans	12.35% - 13.74%

(c) Panel C: Average Rates on Other Loans as of July 2024.

Notes: The tables report various interest rates on different types of (a) federal student loans, (b) mortgage loans, and (c) other private student loans and personal loans. Source: the U.S. Department of Education, Freddie Mac, and SoFi.

lenders.[11] It is not clear how many borrowers are refinancing with private lenders, but we do know that less than 9% of student loan debt is held at private lenders.

3 Data

In this section, I discuss the data used in my analysis. I describe how I construct the relevant measures of consumption and debt, and then briefly discuss the choice of monetary policy shock.

3.1 Credit Panel Data

There is limited data in the United States linking household balance sheet items and consumption. The Consumer Expenditure Survey (CEX), the largest and most commonly used survey of household consumption, has detailed information on consumption and income for households but has limited information on debt and assets. The level of detail required to accurately identify liabilities and precisely classify loan types as required in my setting, is not available in the CEX data.

One alternative is to use detailed data on household debt, available in consumer credit panel data, and construct alternate measures of consumption to capture household responses to monetary policy.

[11] This is not generally a viable option for most households, as private student loans are subject to risk-based pricing and are not collateralized so rates can be significantly higher. Individuals may also lose most of the benefits of the federal student loan program such as relatively generous income repayment plans for certain professions.

Therefore, I use data from the University of California Consumer Credit Panel (UC-CCP), a quarterly database of consumer credit bureau records. The data include two samples, a comprehensive sample of all California records, and a national sample, made up of a random 2% anonymized sample of all U.S. consumers with a bureau record.[12]

Using the UC-CCP data, I can get information on current and starting balances for most loans held by consumers including mortgages, student loans, auto loans, credit cards, and other secured and unsecured loans. There are limited demographic variables but I do observe the age and birth date of individuals in the credit panel, allowing for some household-level controls.

3.1.1 Consumption

There is no direct measure of household consumption in the credit panel data. I proxy for consumption using auto loan origination, following Aruoba et al. (2022),[13] to capture vehicle purchases. The use of vehicle purchases to capture consumption expenditures is not new. Mian et al. (2013) rely on aggregated zip-code level information on new car registrations, along with aggregated credit card transaction data, to overcome the difficulty in finding appropriate data sources with both reliable information on consumption and household balance sheet information. Aruoba et al. (2022) expand on this, using credit panel data and constructing an auto loan origination series to look at the role of household financial constraints in explaining the observed decline in aggregate consumption.

One downside to using auto loan originations to measure vehicle purchases is that it only captures the share of vehicle purchases that are financed; I do not observe cash purchases. In figure 3, I plot my constructed auto loan origination series against the Bureau of Economic Analysis' total vehicle expenditures series over the sample period (2004-2019) to show that the credit bureau data capture aggregate vehicle expenditures. The red line shows the evolution of vehicle purchases as measured in the UC-CCP (normalized to the 2006 value) against the BEA equivalent series, shown in blue. The series align reasonably well over the sample. There are discrepancies following the financial crisis, including a slightly delayed trough in the UC-CCP series, but the general trends match well. The primary concern would be the impact that interest rate changes may have on auto loan financing decisions. Interest rates may be affecting the series in two ways. First, interest rates may change the incentives for financing vehicle purchases instead of paying in cash. Given the results of the above comparison with aggregate sales data (which include cash and financed auto sales), it appears that a change in the percentage of households financing new vehicle purchases is not a significant factor in the aggregate, but could still be a concern at the individual level.[14]

[12]For my analysis, I use a 2% random sample of the national dataset.

[13]Aruoba et al. (2022) use auto loan origination to calculate the average probability of auto loan origination and calculate the dollar value by combining this probability and the average price of a car in the relevant year. I use the origination value of the loan. The key difference is that my series does not include down payments or trade-ins.

[14]Importantly for my setting, any change in financing behavior would only be relevant if it were differentially affected

Figure 3: UC-CCP Auto Loan Origination Series versus BEA Total Vehicle Purchases Series, 2004-2019

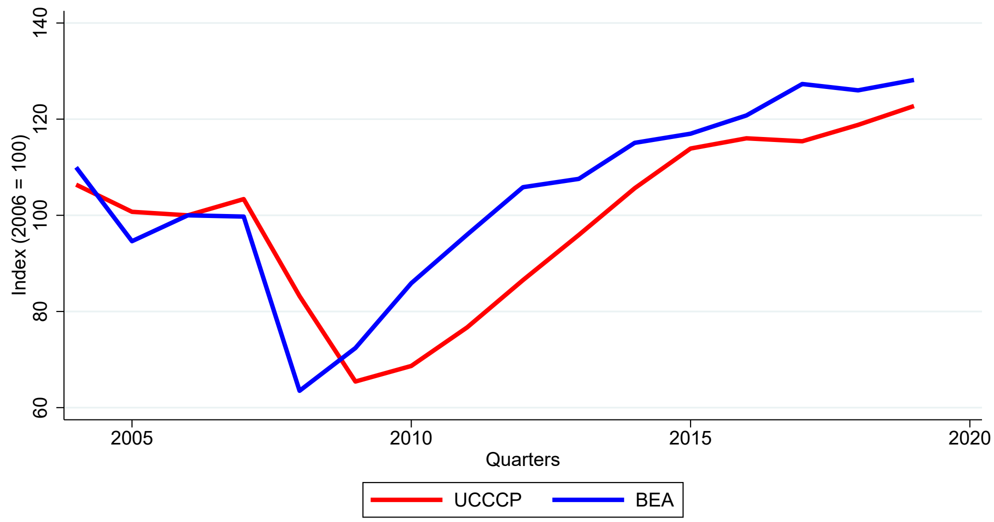

Notes: The figure shows the yearly total consumption expenditures on vehicles from the National Income and Product Accounts (blue), and the auto loan origination series constructed from the University of California, Consumer Credit Panel data (red). Each series is indexed to the value of vehicle expenditures in 2006, shown on the vertical axis. The horizontal axis is years. The BEA series is the sum of new motor vehicles and net purchases of used motor vehicles from NIPA table 2.4.5U, Personal Consumption Expenditures by Type of Product.

The second related concern is auto loan refinancing, changes in interest rates may incentive households to refinance existing loans. In my data this would appear as a new auto loan, but would not map to an actual consumption expenditure. To address this issue, I exclude any new auto loans opened within a 30-day window of the last reported date of an old auto loan closed due to refinancing. I also look at alternate measures of consumption in the UC-CCP, such as credit card balances, to supplement the main specifications.

Now that I have a well-measured variable for vehicle purchases, I can use changes in this loan series to approximate increased spending on cars and other vehicles. Vehicle purchases account for a sizeable share of consumer durable spending. While only a subset of total household spending, consumer durables are particularly responsive to monetary policy (McKay and Wieland 2021), suggesting that, while I am only capturing a portion of consumption spending, I am observing a spending category that

by whether you have fixed-rate or variable-rate student loans or mortgages. The choice to finance is dependent on the current interest rate and an individual's personal credit history which will determine what financing options are available. However, to what extent fixed-rate versus variable-rate debt will factor into this calculation is unclear. Interest rates on existing loans are not included in credit reports and they do not directly influence an individual's credit score, so there should not be differential access to credit based solely the type of interest rate you have on existing debt. The clearest way the change in interest rates may differentially impact the financing decision based on interest rate type is if the increased cash flow from mortgage refinancing or the reduction in variable-rate loan payments causes households who otherwise would have financed their vehicle to purchase the vehicle with cash. In the scenario where an individual has the ability to pay in cash or finance, the choice to finance would likely indicate a preference for liquidity. Potentially, for those with variable-rate or refinanceable debt, the increased cash flow may ease the individual's liquidity constraint, incentivizing them to pay in cash instead of financing. In that case, I would miss that increase in consumption, thereby biasing my estimates downward.

is particularly important to the overall household response to monetary policy. The benefit of this approach is the ability to link vehicle purchases and detailed information on debt balances directly at the borrower level using the UC-CCP data.

Due to data restrictions, I am only able to track individuals in the credit panel data starting in 2004 and to avoid complications from changes in the student loan program during the Covid-19 pandemic, I end my sample in 2019.

3.1.2 Debt Variables

To measure the composition of household balance sheets, I look at the total balances of current loans for each individual. Each unique loan is observed quarterly in the data, and the total quarterly level of debt is simply the sum of current balances for all individual loans by debt type.

For student loan debt, I must define two separate types of debt: fixed and variable rate. Loans opened between July 1992 and June 2006 are classified as variable-rate, and loans opened after July 2006 are classified as fixed-rate.[15] In order to confirm that the credit panel data is capturing the dynamics of the student loan debt in the United States over the last 15 years, I compare the UC-CCP total student loan debt series against the outstanding student loans series from the Federal Reserve's G.19 Consumer Credit statistical release, both indexed to their respective 2006 level.[16] The two series are nearly identical over the time period.

Calculating quarterly household mortgage balances is much simpler. I simply aggregate all loans which are classified as mortgages by quarter for each individual in the primary sample. I similarly plot my UC-CCP total mortgage series against the equivalent series from the Federal Reserve's Financial Accounts statistical release.[17] The calculated mortgage series matches the aggregate trend almost exactly. The UC-CCP series is only slightly higher than the Financial Accounts series but otherwise captures the evolution of mortgage debt.

3.2 Monetary Policy Shocks

In order to determine the effect of monetary policy on household consumption, I must first isolate exogenous unanticipated shocks to interest rates. The measurement of macroeconomic shocks is an ongoing area of research, and there is much debate about the best way to approach measurement. The aim of this paper is not to evaluate individual shock series. One issue with this setting is the timing of the panel dataset. For my purposes, I will use high-frequency shocks from Bauer and Swanson

[15] In Appendix B, I plot the evolution of average balances (left panel) for both borrowers with variable and fixed rate debt and the total level of debt (right panel) for both types. Although variable-rate student loans were phased out in 2006, there is still a significant amount of variable-rate debt throughout the sample.

[16] See appendix B for the figure.

[17] See Appendix B for the figure.

(2023). These shocks measure high-frequency changes in interest rates around FOMC announcements. This approach ensures that the series is plausibly exogenous, as the short window around the FOMC announcements rules out endogeneity issues, like reverse causality, etc. The main benefit of the Bauer and Swanson (2023) shocks is that the time period is well suited to my study.[18]

4 Impact of Monetary Policy By Debt Type

To measure heterogeneous monetary policy responses, I utilize a local projection (LP) framework (Jordà 2005), in which the projected future values of an outcome variable are regressed on current changes in the interest rate, to estimate impulse response functions (IRF). Since monetary policy is conducted in response to current and future expected conditions of the economy, to identify the causal effects of interest rate changes induced by monetary policy, I must use monetary policy shocks to isolate exogenous movements in the real interest rate. Using these shock series, which identify unanticipated changes to the Federal Funds rate, I can estimate the IRF for household-level consumption by estimating a series of simple regressions where the left-hand side variable is consumption projected h horizons in the future, and the right-hand side has indicators for different types of household debt interacted with the monetary policy shocks, along with a set of individual level controls.

The main regression captures the differential responses of different groups of households based on their debt status. Specifically, I estimate the following equation:

$$\begin{aligned} A_{i,t+h} = \sum_{D} \alpha_h^D \delta_i^D + \sum_{l=0}^{L} \sum_{D} \beta_{h,l}^D [\delta_i^D \times \epsilon_{t-l}] + \sum_{l=0}^{L} \gamma_{h,l} \epsilon_{t-l} + \sum_{m=1}^{M} \theta_{h,m} A_{i,t-m} \\ + \sigma_h Age_{i,t} + \sum_{l=0}^{L} \lambda_{h,l} [Age_{i,t} \times \epsilon_{t-l}] + \nu_{i,t} \end{aligned} \tag{1}$$

where $A_{i,t+h}$ is auto purchases of individual, i, h horizons in the future, δ^D is a binary indicating whether household i holds debt of type $D \in$ [fixed rate student loan, variable rate student loan, mortgage]. ϵ_t are the Bauer and Swanson (2023) monetary policy shocks, $Age_{i,t}$ is age of individual i in time t. A key identifying assumption for my results is that there are no additional variables that are driving differential responses to the monetary policy shocks that are correlated with debt status, such as demographic differences between individuals with different debts. I therefore include a set of controls, including age and age interacted with the monetary policy shocks to account for differential monetary policy responses over the life-cycle. I also include lags of the shock and lags of the outcome variable as recommended in Montiel Olea and Plagborg-Møller (2021).[19]

[18] I discuss the choice of shocks further in Appendix F.

[19] In the main specification, I include 16 lags of both the shock and outcome variable. I show in appendix **??** the results

The coefficients of interest are the β_h^Ds, the interaction terms, or the additional consumption of households with debt D over the non-debtor response. These coefficients capture the dynamic effect of monetary policy on household consumption at horizon h; IRFs are created by taking each coefficient over the entire time horizon. Since the shock series are interacted with indicators for debt holding, I can isolate the heterogeneous consumption responses for households with different types of debt.

In my sample, the non-debtor response is the response of individuals with a credit bureau record who do not have mortgage or student loan debt. Ideally, I would observe all individuals regardless of credit history, but in reality, this type of credit panel data will miss individuals who do not have any credit history. There have been efforts to quantify how large this population of "credit invisibles" is and in 2015 the Consumer Financial Protection Bureau estimated that some 26 million U.S. adults had no credit record with any of the three major credit bureaus, representing approximately 11% of the U.S. adult population.[20] So while my results capture a vast majority of U.S. households, there are some individuals who will, by definition, be excluded from the non-debtor responses.

4.1 Mortgagors vs. Fixed Rate Student Loan Holders

In Figure 16, I show the additional consumption response for households with fixed rate student loan debt (top panel) and households with mortgage debt (middle panel) to an expansionary monetary policy shock.[21] These should be interpreted as the additional response of these households relative to the non-debtor response. For those with fixed-rate student loan debt, the response is generally lower than the non-debtor response. This can be seen in the middle panel of Figure 4, which plots the cumulative interaction effect.[22] Here we can see that the fixed rate student loan response is significantly below the non-debtor response until the end of the 16 horizons. Mortgagors on the other hand have a consistently higher response peaking at a cumulative increase in consumption above the non-debtor of around $4000. Importantly, these two responses are statistically different from each other at the 95% level.

The response of mortgagors is consistent with the literature suggesting that mortgage debt is an important channel through which monetary policy is transmitted to household consumption. The comparatively weaker response of individuals with fixed-rate student loan debt is also consistent with the theory that individuals with fixed-rate debt are less responsive to monetary policy.

These results are not a direct test of the importance of interest rate type in the transmission of monetary policy, and one could come up with a number of alternate explanations as to why we may observe these differential responses between mortgagors and student loan holders. The main alternate explanation for why the responses may differ is that mortgagors and student loan holders are

with different lag lengths

[20] https://files.consumerfinance.gov/f/201505_cfpb_data-point-credit-invisibles.pdf

[21] The Federal Funds Rate response is shown in Appendix F.

[22] Cumulative responses are the integral under the impulse response of the level of consumption.

fundamentally different groups and therefore the differential responses are likely driven by differences between groups instead of the interest rates or refinancing mechanism. Another alternate explanation could be that since the average mortgage is larger than the average student loan, the additional response for mortgagors is driven more by the size of the debt, and individuals with larger student loan burdens would likely respond similarly. I show evidence below that neither of these explanations seem to hold in the data, firstly by exploiting the change in interest rate policy in the federal student loan program and comparing the responses of student loan borrowers with fixed rate debt versus those with variable rate student loan debt. I also show alternative specifications which more directly test the importance of the size of debt in generating the differential responses, these results suggest that those with higher student loan balances do not necessarily have higher responses.

4.2 Variable vs. Fixed Rate Student Loan Holders

Again in figure 16, I plot the additional consumption response for households with fixed-rate student loan debt (top panel) and households with variable-rate student loan debt (bottom panel) to an expansionary monetary policy shock. Here we can see that the households with variable rate student loan debt have higher responses relative to those with fixed rate student loan debt, and from Figure 4 we can see that the variable rate student loan response is similar in magnitude to the mortgage response. These results provide a much more direct test of the interest rate mechanism, suggesting that there are important implications for fixed versus variable rate debt on the transmission of monetary policy.

Figure 4: Cumulative Consumption Response to an Expansionary Monetary Policy Shock by Debt Group Relative to Non-Debtor Response

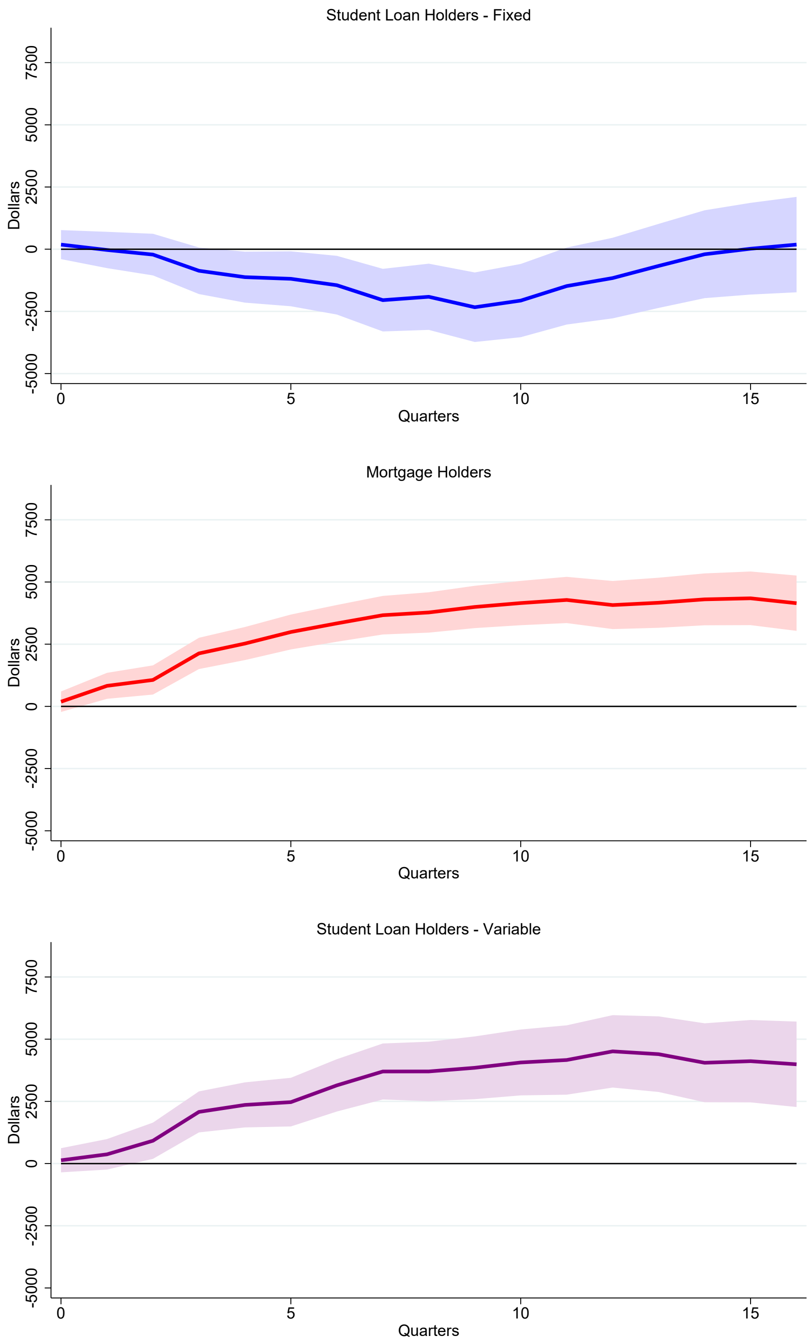

Notes: The figures show the cumulative impulse response of consumption to an expansionary monetary policy shock for households with different types of debt relative to the non-debtor response. Consumption is defined as auto loan originations. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock. Cumulative responses are the integral under the impulse response of the level of consumption.

4.3 Total Household Responses

Given these differential responses, the overarching question is whether these results have broader implications for the conduct of monetary policy. To explore this question further, I look at the total responses for these different households and then explore what the shifting distribution of household debt would imply for average monetary policy responses.

In Figure 5, I plot the total cumulative responses for each type of household in my setting.[23] From Equation 1 this is the sum in each period h, of β_h^D, which captures the additional response to monetary policy driven by holding debt D, and γ_h which captures the effect of the monetary policy shocks which is common to all households. They all have large and significant responses, suggesting strong general equilibrium effects driving monetary policy responses during this period. But, importantly, the mortgagors and variable rate student loan holders, have higher responses than the fixed rate student loan holders.

Figure 5: Total Cumulative Consumption Response to an Expansionary Monetary Policy Shock by Debt Type

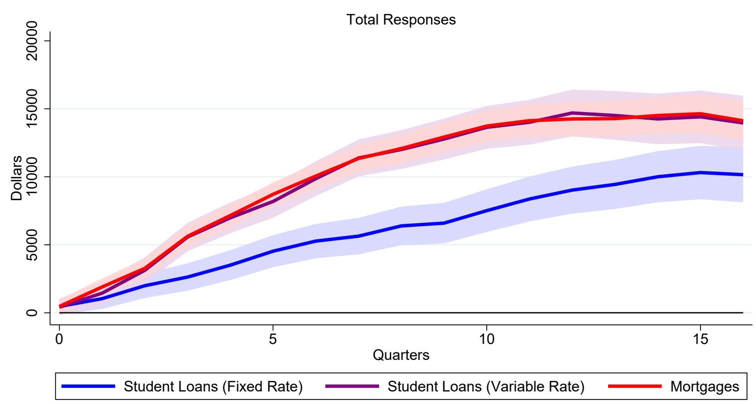

Notes: The figure shows the total cumulative impulse response of consumption to an expansionary monetary policy shock for households with different types of debt. The total response is constructed by summing the non-debtor response and the additional interaction term at each horizon for each of the three types of debt. Consumption is defined as auto loan originations. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents 2007 U.S. dollars, and the horizontal axis is quarters following the shock. Cumulative responses are the integral under the impulse response of the level of consumption.

I then calculate the implied average response given the distribution of household debt in 2007 and 2019. In this calculation, I focus solely on the effect of the shift in the household balance sheet away from refinanceable and variable-rate debt, towards a balance sheet with a larger share of non-refinanceable debt, ignoring other economic factors that may have influenced monetary policy responses and assuming that the general equilibrium effects are constant between the two periods. In equation 2, I take the non-debtor response (γ) and add the additional response for mortgagor and student loan borrowers

[23] Cumulative responses are the integral under the impulse response of the level of consumption.

using the cumulative interaction coefficients (β^D) weighted by the share of households with each type of debt in 2007 and 2019 respectively.

$$\Theta_{year} = \gamma + \delta_{year}^{FixedSl} * \beta^{FixedSL} + \delta_{year}^{VariableSl} * \beta^{VariableSL} + \delta_{year}^{Mortgage} * \beta^{Mortgage} \quad (2)$$

δ_{year}^{D} represents the share of U.S. households holding each type of debt D, from the Survey of Consumer Finances. Taking the difference between the 2007 and 2019 implied responses gives an estimate of how the changing composition of the household balance sheet has affected the transmission of monetary policy to household consumption. In Figure 6, I plot the percent change in the cumulative average household response in each horizon, to show how the changing composition affects the overall response to monetary policy over the four-year horizon.

From this back-of-the-envelope calculation, I estimate that the changing composition of household balance sheets corresponds to, all else equal, a decrease in the average monetary policy response of around 10% after two years. This represents a significant change in the overall monetary policy response, and if trends in student loan and mortgage debt continue on the same path, there could be even greater implications going forward.

Figure 6: Percent Change in the Implied Average Household Response to an Expansionary Monetary Policy Shock at Each Horizon, given the Distribution of Household Debt in 2007 and 2019

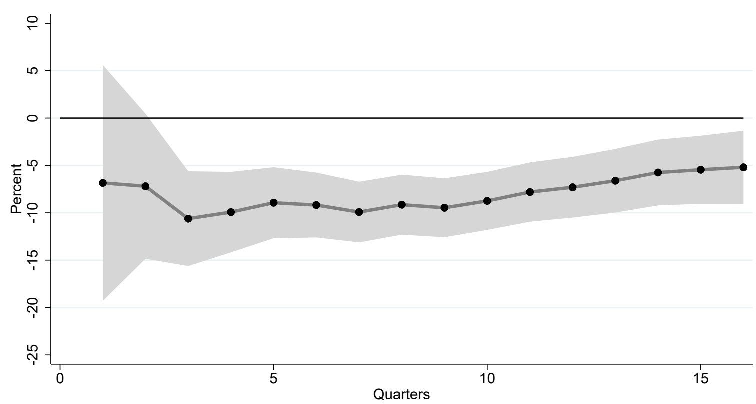

Notes: The figure shows the percent change in the calculated average household response given the distribution of household debt in 2007 and 2019. The implied average response is the non-debtor response (γ) plus the additional response for mortgagor and student loan borrowers using the cumulative interaction coefficients (β^D) weighted by the share of households with each type of debt in 2007 and 2019 respectively. The vertical axis is percent, and the horizontal axis is quarters following the shock.

5 Validating the Mechanism

Below I provide a variety of alternate specifications to further validate the refinancing mechanism and address some of the limitations of the main specification.

5.1 Debt Levels Results

To asses whether the size of the loans is driving the monetary policy response, I estimate an alternate version of the main specification without the binary variable capturing debt types and instead include the level of debt directly. Equation 3 is nearly identical to Equation 1 just with slight modifications. I substitute $X_{i,t}^{D}$ in for δ_i^D, where $X_{i,t}^{D}$ captures the total quarterly balance of household i's debt in category D in time t.

$$
\begin{aligned}
A_{i,t+h} = \sum_{D} \alpha_h^D X_{i,t}^D + \sum_{l=0}^{L} \sum_{D} \beta_{h,l}^D [X_{i,t}^D \times \epsilon_{t-l}] + \sum_{l=0}^{L} \gamma_{h,l} \epsilon_{t-l} + \sum_{m=1}^{M} \theta_{h,m} C_{i,t-m} \\
+ \sigma_h Age_{i,t} + \sum_{l=0}^{L} \lambda_{h,l} [Age_{i,t} \times \epsilon_{t-l}] + \nu_{i,t}
\end{aligned}
\tag{3}
$$

In Appendix D, I show the cumulative interaction effect of households for an additional \$10,000 in fixed-rate student loan debt (top panel), mortgage debt (middle panel), and variable-rate student loan debt (bottom panel). Having an additional \$10,000 in variable-rate student loan debt leads to an increase in consumption following an expansionary shock. By contrast, an additional \$10,000 in fixed-rate student loan debt or mortgage debt does not appear to have significant effects on the household response to consumption.

While the above regression is the most straightforward approach to incorporating debt levels into the regression, the linear nature of the specification allows for large outliers in the data to drive the response. To address this concern, I also estimate a non-parametric version of the specification with households separated into quartile bins based on their debt balances in student loan and mortgage debt. I plot the results of these binned regressions in Figure 7. Here I include the coefficients for each quartile in only select horizons.[24] The consumption response for mortgagors (middle panel) and variable rate student loan holders (bottom panel) is monotonic in the level of debt. This is again consistent with the refinancing mechanism, as higher levels of mortgage debt should increase the incentives to refinance, and would lead to higher reductions in payments when loans are refinanced. Similarly, for variable-rate student loan holders, there should be a direct relationship between the size of the loan and the reduction in payments from changes in rates. For the fixed-rate student loans (top panel), there is not the same monotonic effect.

[24] The full-time horizon plots are included in appendix D.

Figure 7: Cumulative Consumption Response to an Expansionary Monetary Policy Shock by Debt Type Relative to Non-Debtor Response - Quartile Regression, Select Horizons

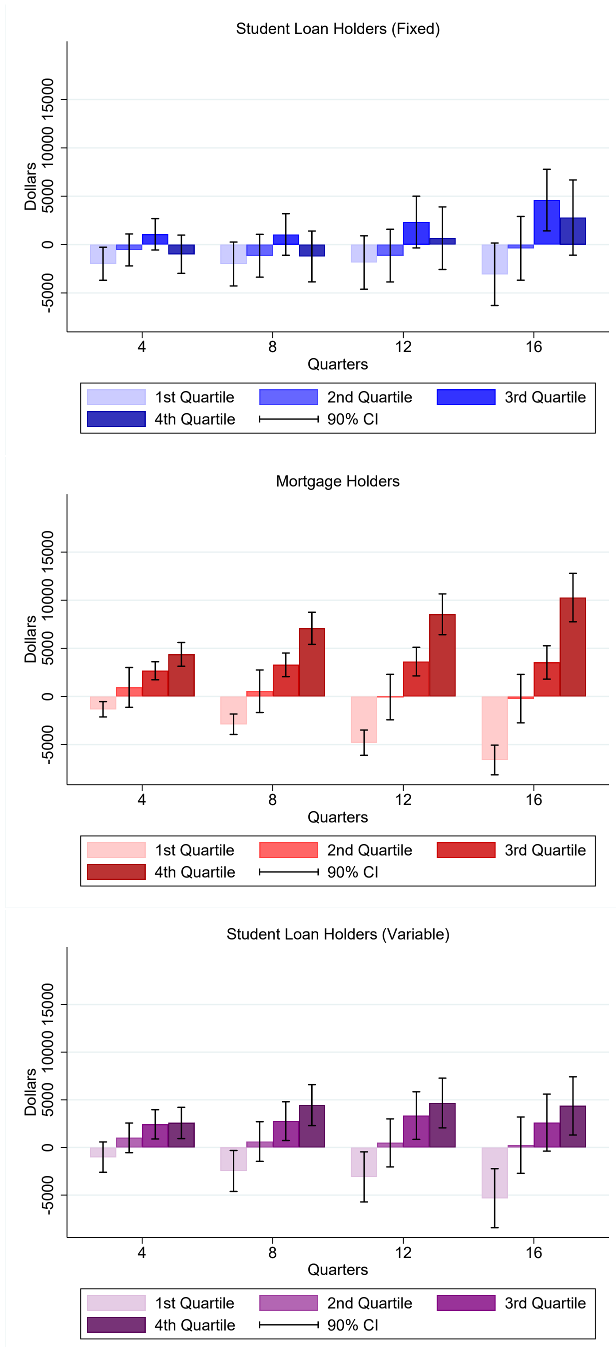

Notes: The figures show the cumulative impulse response of consumption to a one standard deviation expansionary monetary policy shock for households with different levels of debt relative to the non-debtor response. Households are divided into quartile bins according to their level of debt. Consumption is measured as auto loan originations. 90% confidence intervals are shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock. Cumulative responses are the integral under the impulse response of the level of consumption.

5.2 Effect of Household Controls

The goal of this paper is to assess the importance of fixed-rate student loan debt versus refinanceable mortgage debt, and accounting for differences between households with these types of debt is a major challenge. The additional analysis comparing households who got variable-rate student loans before the transition to fixed-rate federal loans in 2006 against households who got student loans after the policy change, acts as a check on the validity of the specification in light of these concerns. The heterogeneous monetary policy response literature often handles this type of issue by controlling for these household demographic characteristics. However, some other key demographic features that can drive monetary policy responses including income, education, and family size are not universally available in the credit bureau data. Limited reliable household controls are a downside of the data, but importantly the differences between households will only cause issues if they cause the households to respond differently to monetary policy shocks.

Looking at the overall composition of mortgagors and student loan holders in the Survey of Consumer Finances (SCF) shows that they are similar along the dimensions which are likely most important for the transmission of monetary policy from the theory, mainly both groups are higher income and more highly educated than the overall population. The key demographic factor where they differ is age. Both mortgagors and student loan holders are younger on average, but the average student loan holder is younger still than the average mortgagor, consequently, the family structure also differs between the groups. Both are factors that likely contribute to any income discrepancies in the data, conditioning on age and marital status in the SCF data brings the two groups closer together (see table 2).

Table 2: SCF Income Comparison - Mortgagors, Student Loan Borrowers, and All Households

Group	Med. total family income	Med. wage income	Med. total family income, couples btw. age 30-40	Med. wage income, couples btw. age 30-40
Mortgage	$81,000	$64,000	$92,000	$87,000
Student Loan	$59,000	$49,000	$85,000	$80,000
All HHS	$50,000	$31,000	$74,000	$66,000

Notes: The table shows the median total family income (column 2) and median wage income (column 3) from the SCF for households with mortgage debt, student loan debt, and for all households. Columns 4 and 5 show the median total family income and median wage income for couples between the ages of 20 and 40 years old. Source: the Federal Reserve Board; Survey of Consumer Finances.

Age accounts for over half of the difference in median household incomes shown above, and given there is reason to believe that monetary policy responsiveness varies over the life-cycle, it is important to control for age. Fortunately, age or birth cohorts can be constructed using the data, I therefore control directly for these life cycle effects in the main specification by including age and age interacted with the monetary policy shocks. To evaluate how important these age controls are for my results, I include a version of the main specification excluding the controls for age and age interacted with the shocks. The results without age controls are nearly identical to the main results. The mortgage response is almost unchanged, and while both the fixed-rate and variable-rate student loan responses are both marginally higher, the mortgage response and variable-rate student remain significantly larger at the 95% confidence level. Suggesting that the differential responses I observe in the main results are not driven by demographic differences across households.

5.3 Evidence of Refinancing

While the effect of monetary policy on refinancing is well established (Wong 2021; Berger et al. 2021; Beraja et al. 2019; Eichenbaum et al. 2022), to further support the interest rate mechanism I provide some evidence that mortgage refinancing increases in response to the monetary policy shocks in the credit panel data. The scope of the data is more limited in this setting but for mortgage loans, I have the necessary information to estimate directly the impact of monetary policy shocks on refinancing and mortgage payments.

5.3.1 Mortgage Rate Results

First, I estimate the response of mortgage rates to the monetary policy shocks. Specifically, I estimate the following regression,

$$\Delta R_t = \alpha + \beta mps_t + \sum_{l=1}^{L} \gamma_{h,l} R_{t-l} + \nu_t. \tag{4}$$

Where ΔR_t is the change in the weekly mortgage rates from Freddie Mac, retrieved from the Federal Reserve Bank of St. Louis. I include 30-year and 15-year fixed-rate mortgages, along with 5/1-year adjustable-rate mortgages. mps_t is the monetary policy surprise as measured by Bauer and Swanson (2023). I also include lags of the interest rate, R_t, and ν_t is the residual.

The β coefficients show the two-week change in the mortgage rate from a one standard deviation monetary policy shock, following Wong (2021).[25] I find that there is a positive relationship between the monetary policy shocks and mortgage rates over, meaning an expansionary monetary policy shock will cause mortgage rates fall.

[25] I also include the change in the daily Market Yield on U.S. Treasury Securities at 10-Year Constant Maturity retrieved from FRED, Federal Reserve Bank of St. Louis.

Table 3: Mortgage Rate Responses to Monetary Policy Shock

Interest Rate	30-Yr FRM	15-Yr FRM	5/1-Yr ARM	10-Yr Treasury
MPS	0.431*	0.677***	0.851***	0.404***
	(.300)	(.301)	(.226)	(.1817784)

Notes: The table shows the response of mortgage rates and the 10-year treasury yield to a one standard deviation monetary policy shock. Standard errors in parentheses.
* $p<0.32$, ** $p<0.10$, *** $p<0.05$

5.3.2 Refinancing

Mortgage loans in the United States tend to be fixed rate but with an option to refinance.[26] As interest rates fall households have an incentive to refinance their existing mortgage debt to reduce payments. In the credit panel data, there are indicators for refinancing, so I can directly measure the effect of the monetary policy shock on refinancing behavior. For the next results, I estimate the following regression on a limited sample including only individuals with mortgage debt.

$$A_{i,t+h} = \alpha + \sum_{l=0}^{16} \gamma_{h,l} \cdot \epsilon_{t-l} + \sum_{m=1}^{16} \theta_{h,m} \cdot A_{i,t-m} + \nu_{i,t} \tag{5}$$

In the refinancing version of the specification, $A_{i,t+h}$ is simply an indicator for whether a household refinanced in quarter $t+h$. The right-hand side includes the monetary policy shocks (ϵ) and lags of the shock and outcome variable. I find that a one standard deviation monetary policy shock leads to a two percentage points increase in mortgage refinancing over two years.

5.3.3 Mortgage Payments

Additionally, estimating the above specification with mortgage payments on the right-hand side, suggests that the monetary policy shocks leads to a cumulative decrease in mortgage payments of $3000 over the eight horizons. The total reduction in payments on mortgage loans over the two shown here seems to match the cumulative increase in consumption after eight quarters from the main results. The similarity in magnitudes suggests that the additional consumption response for individuals with mortgages may be driven by the reduction in payments on mortgage loans.

[26] During the subprime mortgage crisis in the lead-up to the Great Recession, there was also a substantial amount of variable rate mortgage debt. I cannot identify the type of loan in my data, but the inclusion of variable-rate mortgage debt does not impact the interpretation of my results.

Figure 8: Cumulative Response of the Mortgage Refinancing Rate to an Expansionary Monetary Policy Shock

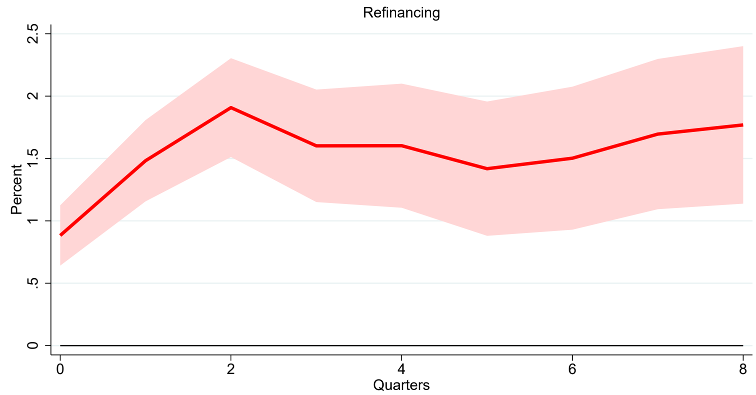

Notes: The figure shows the cumulative impulse response of the mortgage refinancing rate to an expansionary monetary policy shock for households with mortgage debt. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents percentage points, and the horizontal axis is quarters following the shock.

Figure 9: Cumulative Response of Mortgage Payments to an Expansionary Monetary Policy Shock

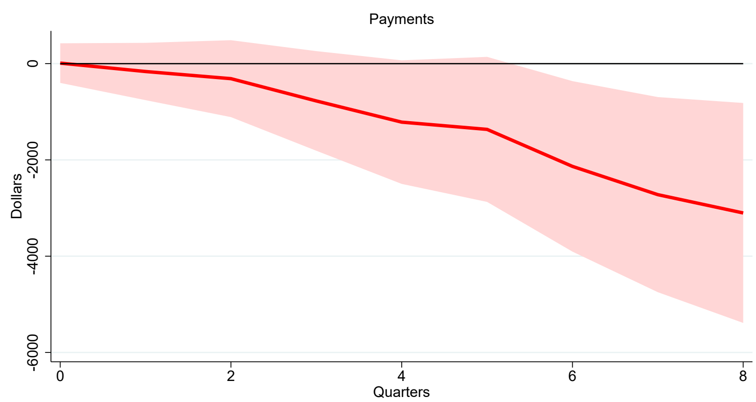

Notes: The figure shows the cumulative impulse response of mortgage payments to an expansionary monetary policy shock for households with mortgage debt. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents dollars, and the horizontal axis is quarters following the shock.

5.4 Robustness Checks

In Appendix G, I include versions of the main results with different lag lengths to check the sensitivity of the main regression to alternate specifications. The results are broadly similar, and changing the number of lags included does not meaningfully change the response.

Then, in Appendix H, I test the robustness of my results to alternate consumption measures and find that the results do not significantly change. I first use credit card debt and test the changes in balances. Mortgagors and households with fixed-rate student loan debt have slightly lower consumption than the non-debtor when measured using credit cards. Households with variable rate student loan debt have a similar, if not slightly higher, credit card response relative to the non-debtor households. However, measuring consumption with credit card balances in the credit panel data is subject to several measurement concerns. I, therefore, include some potentially relevant additional results, including estimating the same regressions with other loan types, such as other secured and unsecured loans, which are not subject to the same measurement issues as credit card balances. These analyses are consistent with the main results that mortgagors and those with variable-rate student loans are more responsive. Since the auto loan origination series is better measured and is more likely to capture contemporaneous consumption without interference from other factors that might impact the current balance on loans, such as refinancing, changes in payments, etc., I use only the auto loan series in the main specification.

6 Model

The cross-sectional results above provide an estimate for the differential micro-level consumption responses and the back-of-the-envelope calculation attempts to quantify the broader implications by calculating the change in the average effect induced by the shift away from variable-rate or refinanceable debt. However, how these estimates correspond to changes in aggregate demand is less clear, as this approach excludes the general equilibrium dynamics and any other aggregate impacts from the changing household balance sheet. In this section, I lay out a simple New Keynesian model with heterogeneity to show under what assumptions the results from my empirical section can be extended to aggregate monetary policy responses.

I find that under standard assumptions, mainly if we believe that labor market and household income responses are not differentially affected by whether you hold fixed-rate or variable-rate debt, then the aggregate consumption response to monetary policy will be determined by a standard Keynesian cross type multiplier, which is independent of the type of household debt. In this case the back-of-the-envelope calculation is a good approximation of the aggregate effect. In fact, in the stylized version of the model below the back-of-the-envelope calculation exactly captures the aggregate effect.

6.1 Three-Agent New Keynesian Model

To model my empirical setting, I write down a stylized "Three-Agent" New Keynesian model, an extension of the standard two-agent version (Galí et al. 2007; Bilbiie 2008). These models typically assume that some fraction (ϵ) of households are constrained and act as hand-to-mouth consumers while the remaining $(1 - \epsilon)$ households are unconstrained.

The unconstrained households correspond to the non-debtors in my empirical setting. These unconstrained households are standard optimizing agents who will consume out of their lifetime income, borrowing and saving over the life-cycle to smooth consumption. Variables for these households are denoted with a superscript s for "savers." The constrained agents in my setting are the borrowers. In the model, these households are at their borrowing constraint because of their existing debt and, therefore, do not have access to the credit market. These households then act as hand-to-mouth consumers, consuming all of their income in each period.

I alter the standard TANK model by splitting the constrained households into two categories: λ fraction of fixed-rate borrowers and $(1-\lambda)$ fraction of variable-rate borrowers. The fixed-rate borrowers in my setting correspond to households with fixed-rate student loan debt, and the variable-rate borrowers correspond to households with mortgage and variable-rate student loan debt.[27] Variables for these households are denoted using superscript f and v, respectively.

6.2 Production

Going forward, I will simplify the model to a two-period version. While I lose some generality, the intuition for my empirical setting remains in the simplified version, and the model can be easily extended to an infinite horizon version.

The production side of the model is mostly standard, but instead of the usual fixed price assumption, I assume fixed wages and flexible prices. In the derivation of the equilibrium, I will impose perfectly rigid wages in period one.

The firm produces a single good using labor, N_t, as their only input.

$$Y_t = N_t. \tag{6}$$

Each household provides labor to the firm, and total labor is the sum of labor for each of the households,

$$N_t = N_t^s + N_t^v + N_t^f. \tag{7}$$

Firms do not differentiate between agents. I assume that total labor for each household type is equal to that type's share of the total population of households. Total labor can then be rewritten as:

$$N_t = (1-\epsilon)N_t + \epsilon(1-\lambda)N_t + \epsilon\lambda N_t. \tag{8}$$

The firm chooses its optimal price given the market wage. In the first period, I assume perfectly

[27] This is a simplification as mortgage debt is generally fixed with an option to refinance. The intent here is not to model the refinancing decision, so I present instead a case where the refinancing rate is equal to 1.

rigid wages, so employment is fully demand-determined. In this specific case the wage setting condition simply becomes $W_1 = W_0$ for all households.

From the firm's first order condition, the labor demand curve is,

$$W_t = P_t. \tag{9}$$

Since wages are not changing, the firm's profit-maximizing price in the first period is equal to $W_1 = W_0$ and thus $P_1 = P_0$.

6.3 Households

Total consumption in the model is simply the sum of consumption for the three agents.

$$C_t = C_t^s + C_t^v + C_t^f \tag{10}$$

6.3.1 Saver Households

The saver households in my model are the same as in the standard New Keynesian model, with the following utility function,

$$U(C_1^s, C_2^s, N_1^s, N_2^s) = \left[\log(C_1^s) - \theta^s \frac{(N_1^s)^{1+\chi}}{1+\chi}\right] + \beta\left[\log(C_2^s) - \theta^s \frac{(N_2^s)^{1+\chi}}{1+\chi}\right] \tag{11}$$

Savers in this model hold the debt of the borrowers and receive interest payments along with their wage income. Unconstrained households are able to borrow and save in the form of bonds with a return equal to the nominal policy rate.[28] The resulting budget constraints are,

$$BC_1 : P_1 C_1^s + B_1 = W_1 N_1^s + i_0 D^v + \bar{i} D^f$$
$$BC_2 : P_2 C_2^s = W_2 N_2^s + (1 + i_1) B_1 + D^v + D^f$$

The unconstrained households maximize total lifetime utility subject to their intertemporal budget constraint, with the standard first-order conditions for optimizing households. The two-period Euler equation is:

$$P_2 C_2^s = \beta(1 + i_1) C_1^s P_1. \tag{12}$$

In the first period, when wages are fixed, households adjust labor to meet current demand. In the

[28] $B_1 = 0$ in equilibrium.

second period, households are back on their labor-supply curve:

$$\theta^s C_2^s (N_2^s)^\chi = \frac{W_2}{P_2}. \tag{13}$$

Unconstrained households choose consumption according to their Euler equation, resulting in standard permanent income hypothesis consumption based on the discounted present value of lifetime wealth.

$$P_1 C_1^s = \frac{1}{(1+\beta)}\left[W_1 N_1^s + i_0 D^v + \bar{i} D^f + \frac{W_2 N_2^s + D^v + D^f}{(1+i_1)}\right] \tag{14}$$

$$P_2 C_2^s = \frac{\beta(1+i_1)}{(1+\beta)}\left[W_1 N_1^s + i_0 D^v + i^f D^f] + \frac{W_2 N_2^s + D^v + D^f}{(1+i_1)}\right]. \tag{15}$$

6.3.2 Variable-Rate Debt

The utility function for variable-rate borrowers is:

$$U(C_1^v, C_2^v, N_1^v, N_2^v) = \left[\log(C_1^v) - \theta^v \frac{(N_1^v)^{1+\chi}}{1+\chi}\right] + \beta\left[\log(C_2^v) - \theta^v \frac{(N_2^v)^{1+\chi}}{1+\chi}\right] \tag{16}$$

Variable-rate borrowers are constrained and, therefore, consume their entire disposable income in each period. These households do not have Euler equations, but they do have the same labor supply conditions as the unconstrained households. In the first period, labor is completely demand determined, but their second-period labor supply curve is,

$$\theta^v C_2^v (N_2^v)^\chi = \frac{W_2}{P_2}. \tag{17}$$

The main feature of the model is the existing exogenous debt D held by borrowers for which they must pay an interest payment in the first period. They must then repay the loan in its entirety in period two. I define $D^v = (1-\lambda)D$, so that the level of variable-rate debt is a fixed portion of exogenously determined total debt.

Below are the nominal budget constraints for these households. Variable-rate borrowers work and receive wage income, $W_t N_t^v$, and in the first period pay interest on their debt equal to $i_0 D^v$. In the second period, borrowers must pay back the loan in full. The interest rate on variable-rate loans is exactly equal to the nominal policy rate in the previous period[29] and will therefore adjust in response

[29] This is an assumption made to simplify the algebra and the final expressions, but can be rationalized if we believe there is some time between the change in the policy rate and the updating of the variable interest rate faced by borrowers.

to changes in the lagged nominal interest rate.

$$P_1 C_1^v = W_1 N_1^v - i_0 D^v \tag{18}$$

$$P_2 C_2^v = W_2 N_2^v - D^v \tag{19}$$

6.3.3 Fixed Rate Debt

The fixed-rate borrowers are identical to the variable-rate borrowers with utility function:

$$U(C_1^f, C_2^f, N_1^f, N_2^f) = \left[\log(C_1^f) - \theta^f \frac{(N_1^f)^{1+\chi}}{1+\chi}\right] + \beta\left[\log(C_2^f) - \theta^f \frac{(N_2^f)^{1+\chi}}{1+\chi}\right] \tag{20}$$

and with a labor supply curve equal to,

$$\theta^f C_2^f (N_2^f)^{\chi} = \frac{W_2}{P_2}. \tag{21}$$

Consumption in each period for the fixed-rate borrowers is:

$$P_1 C_1^f = W_1 N_1^f - \bar{i} D^f \tag{22}$$

$$P_2 C_2^f = W_2 N_2^f - D^f \tag{23}$$

where $D^f = \lambda D$. The main difference between fixed-rate and variable-rate borrowers is that the interest rate for fixed-rate loans is not directly tied to the policy rate and, therefore, will not adjust in response to monetary policy changes.

6.3.4 Monetary Policy Responses

Following a monetary policy shock in the previous period (a change in i_o), the consumption response will work through the budget constraint, affecting the households' cash on hand. I focus on the lagged interest rate in order to abstract away from the standard intertemporal substitution effects. These and other general equilibrium effects will, of course, be present in the empirical setting, but they are not required to capture the key intuition, and therefore, I exclude them from the model.

The partial derivative of first-period consumption with respect to the interest rate for the savers is,

$$\frac{\partial C_1^s}{\partial i_0} = \frac{1}{1+\beta} \frac{1}{P_1} D^v \tag{24}$$

The partial derivatives presented only capture changes in the nominal interest rate. I am not considering how other variables, like income, are changing in general equilibrium. In response to expansionary monetary policy, saver households will receive a slight reduction in their wealth from the reduction in interest payments on the variable-rate debt that they own. Since optimizing households would like to smooth consumption over the life-cycle, their demand decreases by the change in payments times their marginal propensity to consume out of the temporary change in income. The reduction in consumption should then be less than the reduction in interest payments.

For variable-rate borrowers, the partial derivative of first-period consumption with respect to the lagged interest rate is

$$\frac{\partial C_1^v}{\partial i_0} = -\frac{1}{P_1} D^v. \tag{25}$$

In response to expansionary monetary policy, demand for these households will increase by exactly the decrease in interest payments in the first period.

Finally, for the fixed-rate borrowers, the partial derivative of first-period consumption with respect to the lagged interest rate for households with fixed-rate debt is,

$$\frac{\partial C_1^f}{\partial i_0} = 0. \tag{26}$$

These households do not respond to changes in interest rates.

Relating the empirical results from my main specification (Equation 1) to the model, the micro-level demand responses from the model correspond most closely with the total household responses shown in Figure 5. The partial derivatives from the model can be linked to the estimated household responses by the following expressions:

$$\begin{aligned}
&\frac{\partial C_1^s}{\partial i_0} = \gamma \\
&\frac{\partial C_1^v}{\partial i_0} = \beta^{VariableSL} + \gamma \text{ or } \beta^{Mortgage} + \gamma \\
&\frac{\partial C_1^f}{\partial i_0} = \beta^{FixedSL} + \gamma,
\end{aligned}$$

where γ captures the omitted category (non-debtor) response, or the effect of the monetary policy shocks which is common to all households, and β^D captures the additional response to monetary policy driven by holding debt D. As discussed in section 4, all of the households have large and significant responses, suggesting strong general equilibrium effects, but the model abstracts from most of the GE effects. Therefore, the interaction term maps most closely with the micro responses in the model with the "saver" households as the omitted category, and the relative treatment effect captured by the interaction term, β^D in Equation 1, represents the additional response induced by holding each type of debt:

$$\frac{\partial C_1^v}{\partial i_0} - \frac{\partial C_1^s}{\partial i_0} = \beta^V \text{or} \beta^M$$
$$\frac{\partial C_1^f}{\partial i_0} - \frac{\partial C_1^s}{\partial i_0} = \beta^F.$$

6.4 Micro-level Demand Response

The change in total consumption in response to the monetary policy shock in period zero can be written as the sum of the partial derivatives of individual consumption for each agent.

$$\frac{\partial C_1}{\partial i_0} = \frac{\partial C_1^s}{\partial i_0} + \frac{\partial C_1^v}{\partial i_0} + \frac{\partial C_1^f}{\partial i_0} \tag{27}$$

Substituting equations 24, 25, and 26 into the expression for the change in household consumption gives:

$$\frac{\partial C_1}{\partial i_0} = \left[\frac{1}{1+\beta} - 1\right] \frac{D^v}{P_1} \tag{28}$$

where D^v can be expressed as $(1-\lambda)D$, the variable rate share of total debt. The response of demand in this model depends on the debt structure (variable vs. fixed rate) and the marginal propensities to consume (MPCs) of the different households.

In the model, the constrained households have an MPC = 1 by definition. The MPC for the saver household in the model is equal to $\frac{1}{1+\beta}$. In an infinite horizon setting, this converges to $(1-\beta)$ with $0 \leq \beta \leq 1$, the familiar MPC in most representative agent New Keynesian models. When discount rates are calibrated appropriately, the MPC is generally close to zero. To match that in my setting, I can set β sufficiently large, as β can be greater than one in the two-period setting. Using these definitions, equation 28 can be expressed in terms of MPCs,

$$\frac{\partial C_1}{\partial i_0} = \left[MPC^s - MPC^v\right] \frac{D^v}{P_1}. \tag{29}$$

If the MPC for the unconstrained households is equal to one ($\beta = 0$), then the presence of variable-rate debt acts as a redistribution channel between borrowers and savers, and the movement toward fixed-rate debt merely shuts down this channel. If we assume that the MPC for the saver households in the model is less than one, then there will be implications for aggregate demand.

Standard models with rational expectations and no credit frictions predict low MPCs as unconstrained households consume according to their permanent income. There is empirical evidence to support heterogeneity in MPCs across borrowers and savers. Generally, savers will be wealthier, older, and have higher incomes, all of which generally result in lower sensitivity to small changes in temporary income. In my empirical analysis, I find large relative responses for the variable-rate student borrowers, corresponding to,

$$\frac{\partial C_1^v}{\partial i_0} - \frac{\partial C_1^s}{\partial i_0} = -(\text{MPC}^v + \text{MPC}^s)\frac{D^v}{P_1} > 0$$

and little to no significant difference between the fixed-rate borrowers and the non-debtors.

$$\frac{\partial C_1^f}{\partial i_0} - \frac{\partial C_1^s}{\partial i_0} = -(0 + \text{MPC}^s)\frac{D^v}{P_1} \approx 0.$$

In the model, the lack of response for the fixed-rate borrowers is expected as their debt payments are unaffected by monetary policy. If the non-borrower response is similar to that of the fixed-rate borrowers in the model, this would imply that the MPC of the saver households is close to zero. This suggests that the large differential response between the variable-rate borrowers and the savers has implications for aggregate demand as $MPC^s - MPC^v \neq 0$.

To estimate the implications for aggregate monetary policy responses in my empirical setting, I use the back-of-the-envelope calculation (equation 2) to estimate the effect of the changing distribution on aggregate demand. The back-of-the-envelope calculation uses the estimated relative treatment effects, holding fixed the share of non-debtors, weighted by share of households with each type of debt in 2007 versus 2019, to capture the effect of the changing distribution of household debt towards more fixed-rate debt. This is equivalent to increasing λ and thus decreasing the level of D^v present in the model. The question is, then, how good of an approximation is the back-of-the-envelope calculation for the aggregate effect?

There is a large literature on how to take micro estimates of consumption responses and derive macro estimates of aggregate demand. Generally, in these models, the general equilibrium effect would be given by,

$$\frac{\mathrm{d}Y_1}{\mathrm{d}i_0} = \frac{\partial C_1}{\partial i_0} \times \Psi$$

where Ψ is a GE multiplier. The total derivative of Y_1 will capture both the effect of the change in the interest rate on consumption and the general equilibrium effect on income as a result of increased household demand. In Appendix I, I derive Ψ in the model and give a parameterization of the full model to show under what assumptions the back-of-the-envelope calculation fully captures the general equilibrium effects.

Taking the derivative of Y_1 with respect to i_0 results in the expected general equilibrium response,

$$\frac{\mathrm{d}Y_1}{\mathrm{d}i_0} = \underbrace{\frac{1}{(1-\frac{1}{1+\beta})(1-\epsilon)}}_{\text{Multiplier, }\Psi} \underbrace{\left(\frac{1}{1+\beta}-1\right)\frac{D^v}{P_1}}_{\frac{\partial C_1}{\partial i_0}}. \tag{30}$$

In my setting, the back-of-the-envelope calculation is not scaling the household-level responses to construct a level estimate of the aggregate effect but instead is meant to show the change in the aggregate response from the shift from variable-rate to fixed-rate debt. As long as the type of debt held by the borrowers does not directly affect the GE multiplier, then $\frac{\Psi^{New}}{\Psi^{Old}} = 1$, and the change in aggregate demand will be equal to the change in household consumption.

$$\frac{\frac{\mathrm{d}Y^{New}}{\mathrm{d}i_0}}{\frac{\mathrm{d}Y^{Old}}{\mathrm{d}i_0}} - 1 = \frac{\frac{\partial C^{New}}{\partial i_0}}{\frac{\partial C^{Old}}{\partial i_0}} - 1 \tag{31}$$

In Equation 30, we can see that the multiplier does not depend on the composition of fixed versus variable rate debt in the economy. It depends only on the fraction of saver households. Therefore, In this simplified version of the model, the back-of-the-envelope calculation for the change in the average monetary policy response would be a good approximation of the aggregate effect.

There are alternate closures of the model that would generate GE multipliers where the composition of household debt would directly affect aggregate demand. For instance, if the type of debt differentially impacted the household's labor supply curve or if labor productivity was introduced and was dependent on the type of debt. While this might be relevant when thinking about mortgages versus fixed-rate student loan debt, it seems less reasonable when thinking about fixed-rate versus variable-rate federal student loan debt, where the debt type is determined by government policy and not by differences across households. Therefore, the intuition from this simple model can be informative for thinking about the implications of increasing fixed-rate student loan debt on the aggregate monetary policy response.

7 Conclusion

In summary, I find that households with fixed-rate student loan debt are less responsive overall to monetary policy shocks compared to households with mortgage debt, consistent with the theory that the structure of debt and the interest rate exposure of households is a key driver of monetary policy responses. I also find that those with variable-rate student loan debt are more responsive to expansionary monetary policy shocks than those with fixed-rate student loan debt, further suggesting that interest rate exposure is an important mechanism for explaining overall responses to monetary policy shocks.

Why should we care about these differential responses by debt holding? Over the last 20 years, the United States has experienced a meaningful upward trend in the share of student loan debt and a parallel decrease in the share of mortgage debt. If these consumption effects are quantitatively important then, combined with the changing composition of household debt, they imply a dampening of aggregate responses to monetary policy. And if these trends continue then the conduct of monetary policy could look very different going forward.

The rising cost of college and the unprecedented growth in student loan debt have led to many discussions around the future shape of student loan policy. Many of these discussions have centered around individual outcomes at the borrower level, but much less attention has been paid to the macroeconomic implications of increasing student loan debt. Understanding the ways in which student loan debt affects household consumption behavior can help to shed light on the unintended consequences of student loan debt. In the context of monetary policy, mortgage debt, and refinancing play an important role in the transmission of monetary policy to real economic activity. The rise in student loan debt, and the changing composition of household balance sheets, may limit the ability of households to take advantage of favorable interest rates through refinancing thus muting the impact of monetary policy. Differential access to refinancing in response to monetary stimulus can also result in substantial consumption inequality (Beraja et al., 2019), increasing existing inequity and further disadvantaging already marginalized groups who are more likely to hold student loan debt. Through this paper, I hope to provide some new insight into the macroeconomic consequences of student loan debt and provide a new perspective to help inform discussions on student loan policy including interest rates, repayment plans, and debt forgiveness, as well as speaking to broader debates on college affordability and access.

References

SeHyoun Ahn, Greg Kaplan, Benjamin Moll, Thomas Winberry, and Christian Wolf. When Inequality Matters for Macro and Macro Matters for Inequality. *NBER Macroeconomics Annual*, 32:1–75, April 2018. ISSN 0889-3365. doi: 10.1086/696046. URL https://doi.org/10.1086/696046. Publisher: The University of Chicago Press.

Gene Amromin, Janice Eberly, and John Mondragon. The Housing Crisis and the Rise in Student Loans. December 2017.

S. Borağan Aruoba, Ronel Elul, and Sebnem Kalemli-Özcan. Housing Wealth and Consumption: The Role of Heterogeneous Credit Constraints. *National Bureau of Economic Research Working Paper Series*, No. 30591, 2022. doi: 10.3386/w30591. URL http://www.nber.org/papers/w30591.

Adrien Auclert. Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6): 2333–2367, June 2019. ISSN 0002-8282. doi: 10.1257/aer.20160137. URL https://pubs.aeaweb.org/doi/10.1257/aer.20160137.

Michael D. Bauer and Eric T. Swanson. A Reassessment of Monetary Policy Surprises and High-Frequency Identification. *NBER Macroeconomics Annual*, 37(1):87–155, 2023. Publisher: The University of Chicago Press Chicago, IL.

Martin Beraja, Andreas Fuster, Erik Hurst, and Joseph Vavra. Regional Heterogeneity and the Refinancing Channel of Monetary Policy*. *The Quarterly Journal of Economics*, 134(1):109–183, February 2019. ISSN 0033-5533, 1531-4650. doi: 10.1093/qje/qjy021. URL https://academic.oup.com/qje/article/134/1/109/5089981.

David Berger, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra. Mortgage Prepayment and Path-Dependent Effects of Monetary Policy. *American Economic Review*, 111(9):2829–2878, September 2021. ISSN 0002-8282. doi: 10.1257/aer.20181857. URL https://pubs.aeaweb.org/doi/10.1257/aer.20181857.

Florin O. Bilbiie. Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of Economic Theory*, 140(1):162–196, May 2008. ISSN 0022-0531. doi: 10.1016/j.jet.2007.07.008. URL https://www.sciencedirect.com/science/article/pii/S0022053107001159.

Sandra E. Black, Jeffrey T. Denning, Lisa J. Dettling, Sarena Goodman, and Lesley J. Turner. Taking It to the Limit: Effects of Increased Student Loan Availability on Attainment, Earnings, and Financial Well-Being. *American Economic Review*, 113(12):3357–3400, December 2023. ISSN 0002-8282. doi: 10.1257/aer.20210926. URL https://pubs.aeaweb.org/doi/10.1257/aer.20210926.

Chunya Bu, John Rogers, and Wenbin Wu. A Unified Measure of Fed Monetary Policy Shocks. *Finance and Economics Discussion Series*, 2019(043), June 2019. ISSN 19362854. doi: 10.17016/FEDS.2019.043. URL https://www.federalreserve.gov/econres/feds/a-unified-measure-of-fed-monetary-policy-shocks.htm.

James Cloyne, Clodomiro Ferreira, and Paolo Surico. Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism. *The Review of Economic Studies*, 87(1):102–129, January 2020. ISSN 0034-6527, 1467-937X. doi: 10.1093/restud/rdy074. URL https://academic.oup.com/restud/article/87/1/102/5272505.

Daniel H. Cooper and J. C. Wang. Student Loan Debt and Economic Outcomes. *Current Policy Perspectives*, 2014.

Natalie Cox. Evidence from Borrower Repayment Decisions. 2017.

Davide Debortoli and Jordi Galí. Monetary policy with heterogeneous agents: Insights from TANK models. Technical report, Department of Economics and Business, Universitat Pompeu Fabra, 2017. URL https://ideas.repec.org/p/upf/upfgen/1686.html. Issue: 1686.

Jason Delisle. Federal student loan interest rates: history, subsidies, and cost. February 2012.

Martin Eichenbaum, Sergio Rebelo, and Arlene Wong. State-Dependent Effects of Monetary Policy: The Refinancing Channel. *American Economic Review*, 112(3):721–761, March 2022. ISSN 0002-8282. doi: 10.1257/aer.20191244. URL https://pubs.aeaweb.org/doi/10.1257/aer.20191244.

Martin Flodén, Matilda Kilström, Jósef Sigurdsson, and Roine Vestman. Household Debt and Monetary Policy: Revealing the Cash-Flow Channel. *The Economic Journal*, 131(636):1742–1771, May 2021. ISSN 0013-0133, 1468-0297. doi: 10.1093/ej/ueaa135. URL https://academic.oup.com/ej/article/131/636/1742/6039341.

Scott L. Fulford and Scott Schuh. Revolving versus Convenience Use of Credit Cards: Evidence from U.S. Credit Bureau Data. *Journal of Money, Credit and Banking*, 55(7):1667–1701, 2023. ISSN 1538-4616. doi: 10.1111/jmcb.13023. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jmcb.13023. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jmcb.13023.

Jordi Galí, J. David López-Salido, and Javier Vallés. Understanding the Effects of Government Spending on Consumption. *Journal of the European Economic Association*, 5(1):227–270, March 2007. ISSN 1542-4766. doi: 10.1162/JEEA.2007.5.1.227. URL https://doi.org/10.1162/JEEA.2007.5.1.227.

Jason N. Houle and Lawrence Berger. Is Student Loan Debt Discouraging Homeownership among Young Adults? *Social Service Review*, 89(4):589–621, December 2015. ISSN 0037-7961. doi: 10.

1086/684587. URL https://www.journals.uchicago.edu/doi/full/10.1086/684587. Publisher: The University of Chicago Press.

Tullio Jappelli and Annalisa Scognamiglio. Interest rate changes, mortgages, and consumption: evidence from Italy. *Economic Policy*, 33(94):183–224, April 2018. ISSN 0266-4658. doi: 10.1093/epolic/eiy001. URL https://doi.org/10.1093/epolic/eiy001.

Òscar Jordà. Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182, February 2005. ISSN 0002-8282. doi: 10.1257/0002828053828518. URL https://pubs.aeaweb.org/doi/10.1257/0002828053828518.

Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743, March 2018. ISSN 0002-8282. doi: 10.1257/aer.20160042. URL https://pubs.aeaweb.org/doi/10.1257/aer.20160042.

Ralph Luetticke. Transmission of Monetary Policy with Heterogeneity in Household Portfolios. *American Economic Journal: Macroeconomics*, 13(2):1–25, April 2021. ISSN 1945-7707, 1945-7715. doi: 10.1257/mac.20190064. URL https://pubs.aeaweb.org/doi/10.1257/mac.20190064.

Alisdair McKay and Johannes F. Wieland. Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy. *Econometrica*, 89(6):2717–2749, 2021. ISSN 0012-9682. doi: 10.3982/ECTA18821. URL https://www.econometricsociety.org/doi/10.3982/ECTA18821.

Alvaro Mezza, Daniel Ringo, Shane Sherlund, and Kamila Sommer. Student Loans and Homeownership. *Journal of Labor Economics*, 38(1), 2020. doi: 10.1086/704609. URL https://www.journals.uchicago.edu/doi/epdf/10.1086/704609.

Atif Mian, Kamalesh Rao, and Amir Sufi. Household Balance Sheets, Consumption, and the Economic Slump*. *The Quarterly Journal of Economics*, 128(4):1687–1726, November 2013. ISSN 0033-5533. doi: 10.1093/qje/qjt020. URL https://doi.org/10.1093/qje/qjt020.

José Luis Montiel Olea and Mikkel Plagborg-Møller. Local Projection Inference Is Simpler and More Robust Than You Think. *Econometrica*, 89(4):1789–1823, 2021. ISSN 1468-0262. doi: 10.3982/ECTA18756. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA18756. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA18756.

Jiri Slacalek, Oreste Tristani, and Giovanni L. Violante. Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area. *Journal of Economic Dynamics and Control*, 115:103879, 2020. ISSN 0165-1889. doi: https://doi.org/10.1016/j.jedc.2020.103879. URL https://www.sciencedirect.com/science/article/pii/S0165188920300488.

Arlene Wong. Refinancing and The Transmission of Monetary Policy to Consumption. March 2021.

A Appendix: Debt Composition Detail

Figure 10: Composition of Total Household Debt, 2007-2020

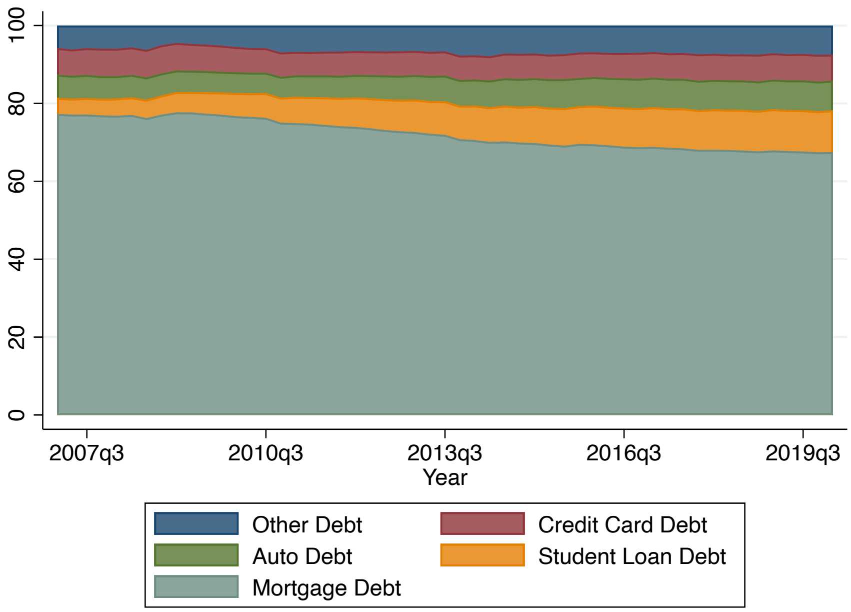

Notes: This figure shows the change in the composition of household debt from 2007 through 2019. Each segment represents the cumulative share of total debt starting from largest debt category to smallest: mortgage debt, student loan debt, auto debt, credit card debt, and all other debt respectively. The vertical axis is percent and the horizontal axis is quarters between 2007q1 and 2019q4. Source: Federal Reserve Board; Z.1 Financial Accounts of the United States and the G.19 Consumer Credit Release.

Figure 11: Student Loan and Mortgage Debt in the Survey of Consumer Finances (SCF), 1989-2019

(a) Share of Households with Student Loan and Mortgage Debt

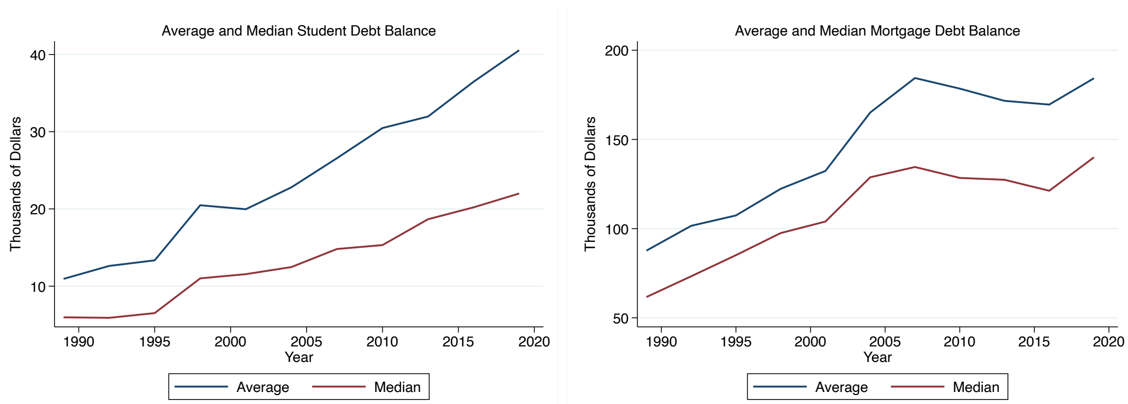

(b) Average and Median Balances for Borrowers with Student Loan and Mortgage Debt

Notes: (a) Figures show the percentage of U.S. households with student loan debt (left panel) and mortgage debt (right panel) between 1989 and 2019. The vertical axis is percent and the horizontal axis is years. (b) Figures show the average (blue) and median (red) balances for households with student loans (left panel) and households with mortgage loans (right panel). Source: Federal Reserve Board, Survey of Consumer Finances. The observations are triennial SCF waves.

Figure 12: Auto Loan and Credit Card Debt as a Share of Total Household Debt, 2007-2019

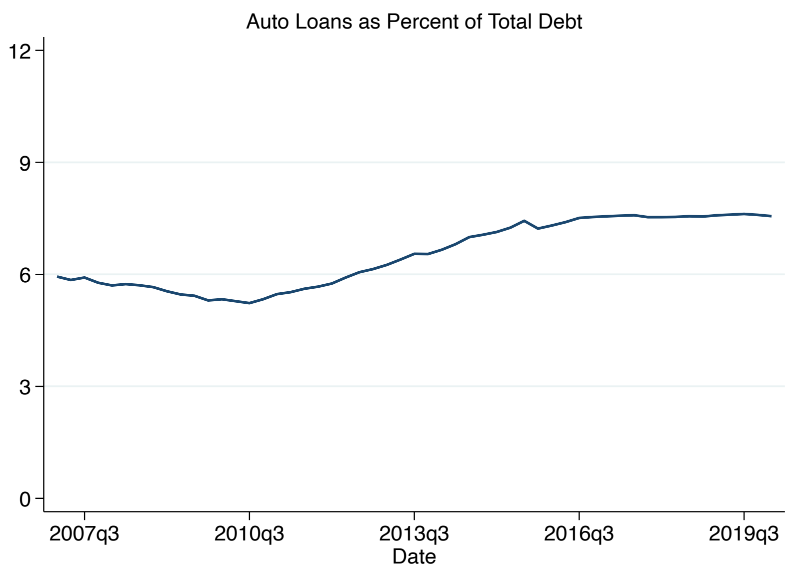

(a) Share of Auto Loan Debt

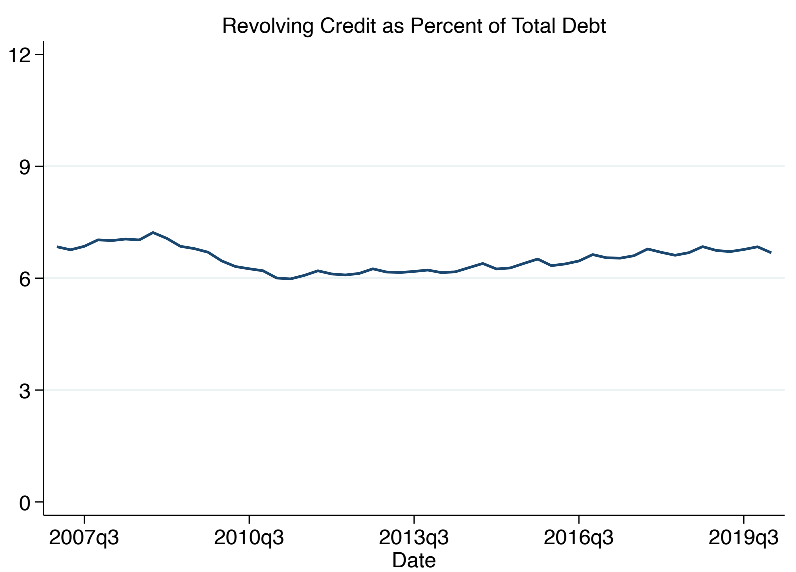

(b) Share of Credit Card Debt

Notes: The figures show the quarterly share of total household debt that is auto loan loan debt (top panel) and credit card debt (bottom panel). The vertical axis is percent and the horizontal axis is quarters between 2007q1 and 2019q4. Source: Federal Reserve Board; Z.1 Financial Accounts of the United States, and G.19 Consumer Credit Release.

B Appendix: Household Debt in the UC-CCP

Figure 13: Fixed-Rate and Variable-Rate Student Loan Debt in the UC-CCP, 2004-2019

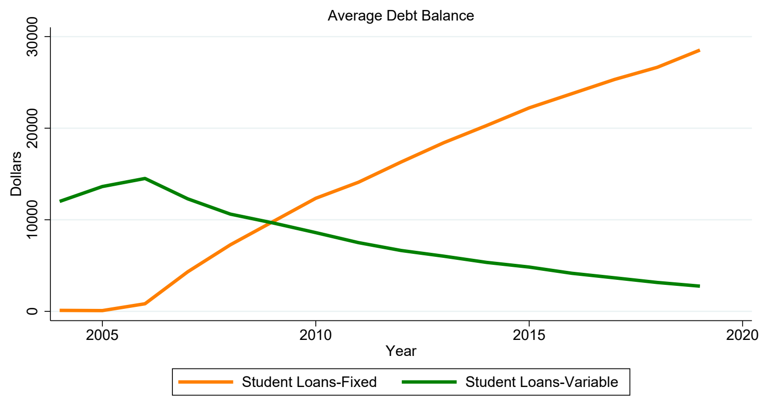

(a) Average Balances for Fixed-Rate and Variable-Rate Student Loan Borrowers

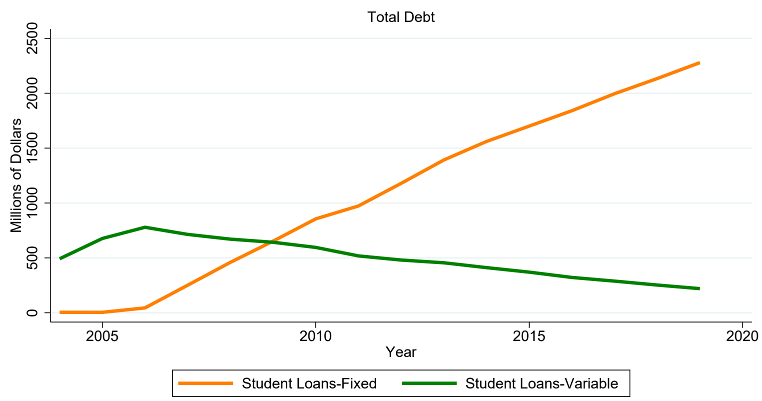

(b) Total Fixed-Rate and Variable-Rate Student Loan Debt

Notes: (1) The figure shows the average fixed-rate student loan balance (orange) for borrowers and variable-rate student loan balance (green) for borrowers in the UC-CCP sample. The vertical axis represents 2006 U.S. dollars. (2) Figure shows the value of total debt in the UC-CCP sample for fixed-rate student loan debt (orange) and variable-rate student loan debt. The vertical axis represents millions of 2006 U.S. dollars.

Figure 14: UC-CCP Student Loan Series and G.19 Student Loan Series

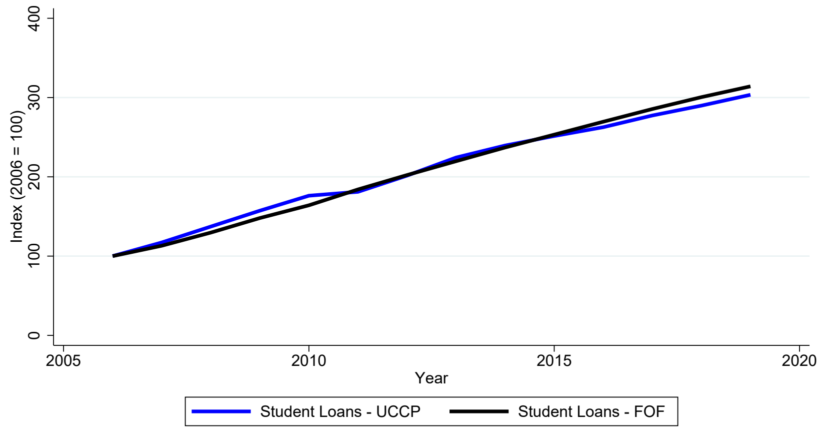

Notes: This figure shows the growth of student loan debt in the UC-CCP (blue) and the aggregate measure of student loan debt from the Federal Reserve Board (black). Each series is indexed to the value of vehicle expenditures in 2006, shown on the vertical axis. The horizontal axis is years.

Figure 15: UC-CCP Mortgage Series and Financial Accounts Mortgage Liability Level Series

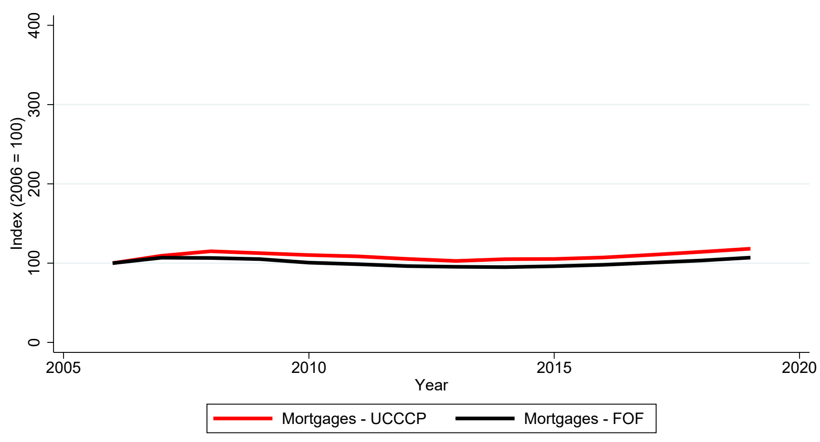

Notes: The figure shows the growth of mortgage debt in the UC-CCP (blue) and the aggregate measure of mortgage debt from the Federal Reserve Board (black). Each series is indexed to the value of vehicle expenditures in 2006, shown on the the vertical axis. The horizontal axis is years.

C Appendix: Additional Results

Figure 16: Consumption Response to an Expansionary Monetary Policy Shock by Debt Group Relative to the Non-Debtor Response

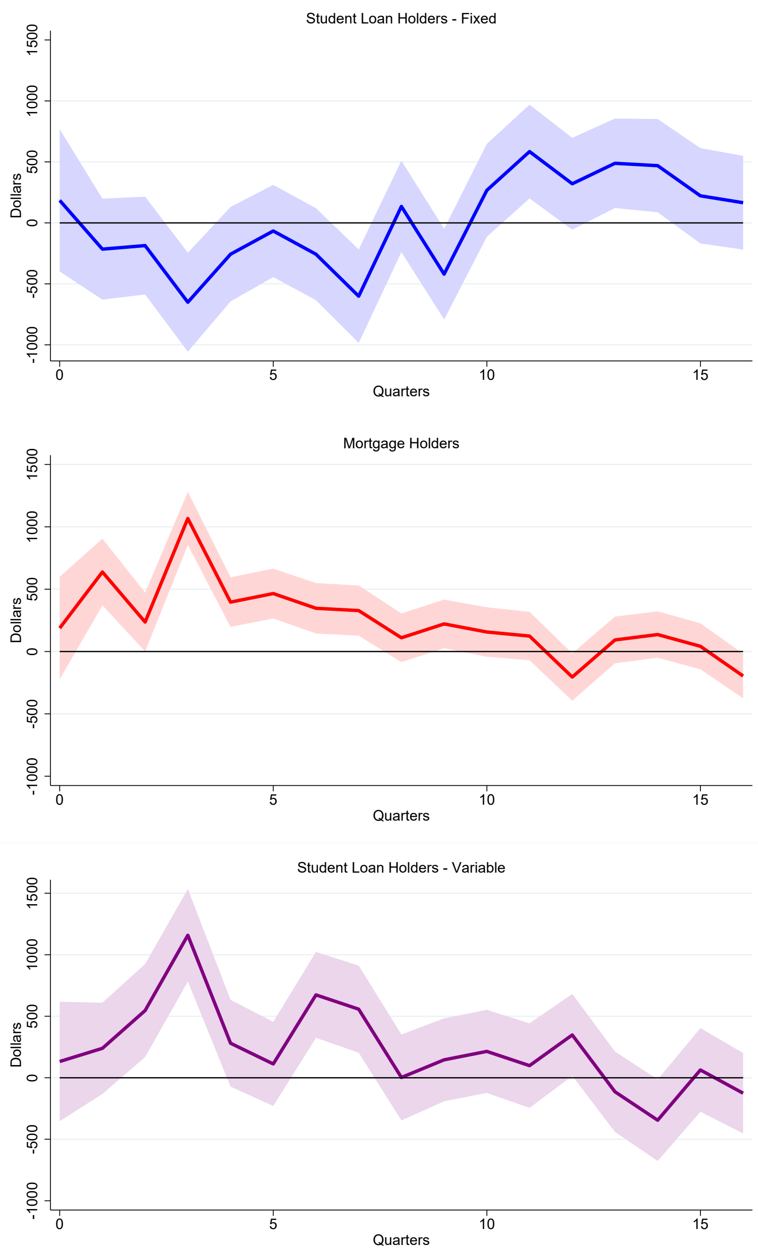

Notes: The figures show the impulse response of consumption to an expansionary monetary policy shock for households with different types of debt relative to the non-debtor response. Consumption is defined as auto loan originations. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock.

Figure 17: Consumption Response to an Expansionary Monetary Policy Shock by Debt Group Relative to Non-Debtor Response - No Age Control

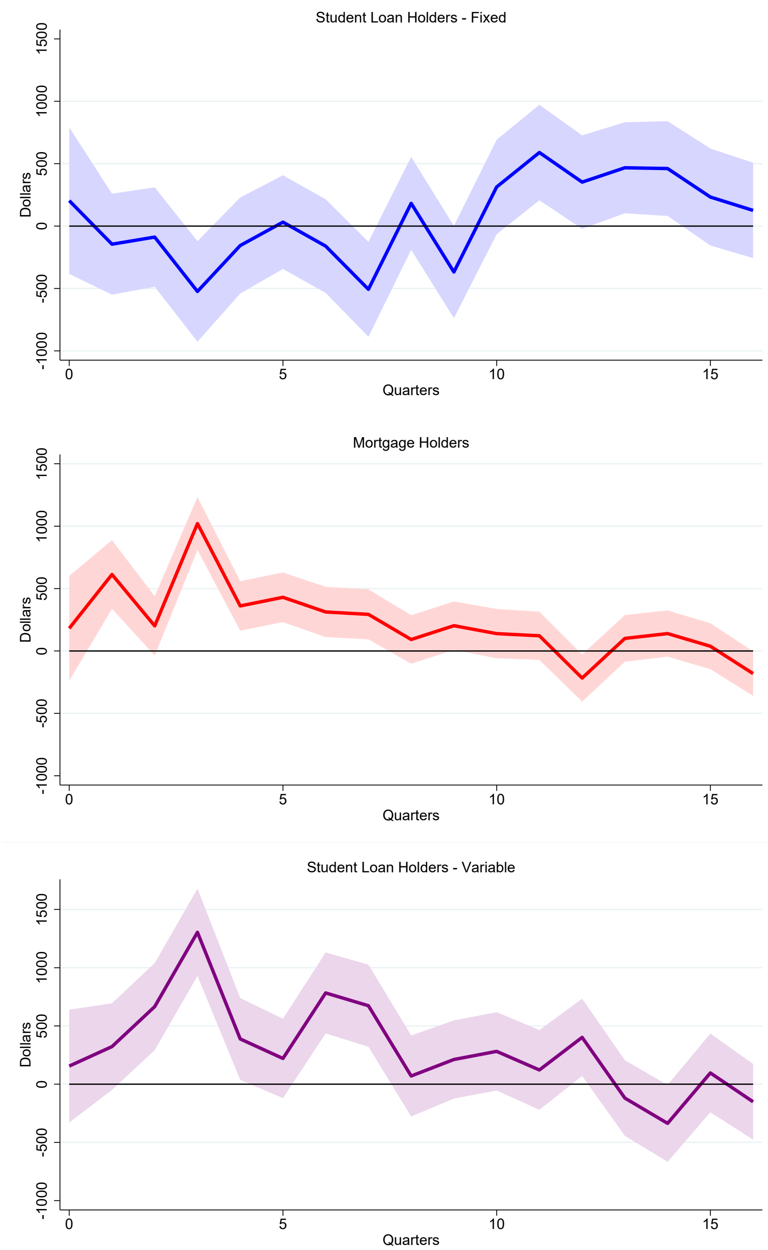

Notes: The figure shows the impulse response of consumption for households with different types of debt relative to the non-debtor response, without the controls for age and age interacted with the monetary policy shock. Consumption is defined as auto loan originations. The shock is a one standard deviation decrease in the monetary policy shock series, which generates a 25 basis point expansionary shock to the Federal Funds Rate. 90% confidence bands shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock.

D Appendix: Level Specification Results

Figure 18: Cumulative Consumption Response to an Expansionary Monetary Policy Shock - Effect of an Additional $10,000 in Debt

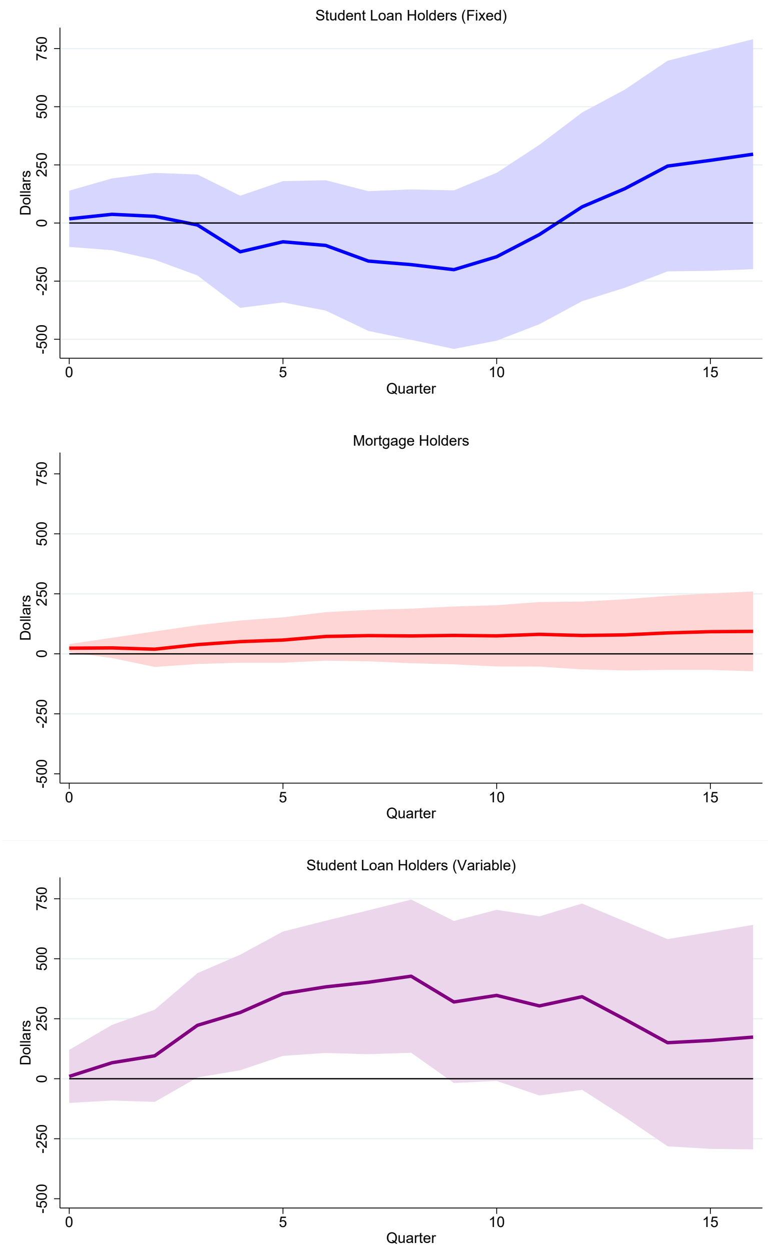

Notes: The figure shows the cumulative impulse response of consumption for an additional $\$10,000$ in each type of debt. Consumption is defined as auto loan originations. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock. Cumulative responses are the integral under the impulse response of the level of consumption.

E Appendix: Quartile Bin Specification Results - All Horizons

Figure 19: Cumulative Consumption Response to an Expansionary Monetary Policy Shock by Debt Type Relative to Non-Debtor Response - Quartile Regression,All Horizons

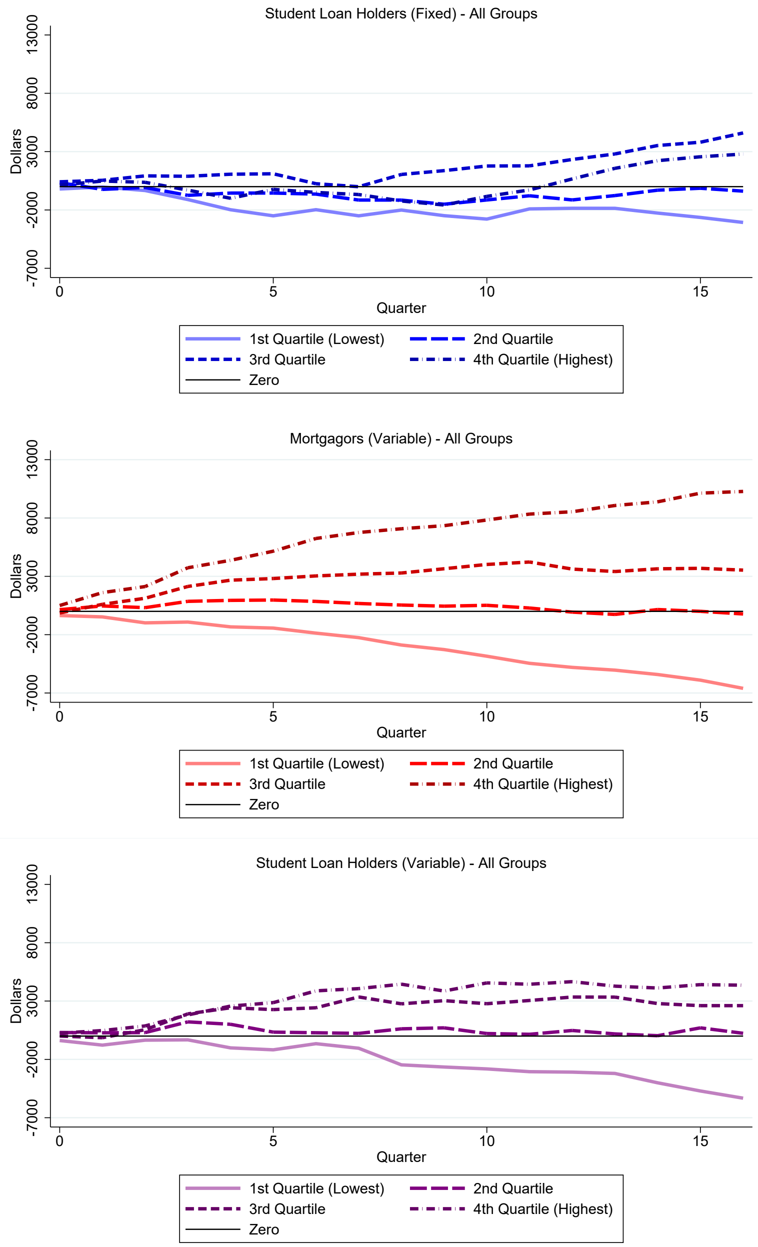

Notes: The figure shows the cumulative impulse response of consumption to a one standard deviation expansionary monetary policy shock for households with different levels of debt relative to the non-debtor response. Households are divided into quartiles bin according to their level of debt. Consumption is measured as auto loan originations. 90% confidence intervals are shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock. Cumulative responses are the integral under the impulse response of the level of consumption.

F Appendix: Monetary Policy Shocks

There are several monetary policy shock series available, but many of the most popular ones do not cover my time period. For this reason, I explore two newer shock series that are well suited to the sample period, the Bauer and Swanson (2023) shocks (BSO) and the Bu et al. (2019) shocks (BRW).

To evaluate the monetary policy shock series, I use aggregate time series data to estimate standard local projections. I also test the shock on the UC-CCP auto loan origination series by estimating simple panel-data local projections (the main specification without any interaction terms) to capture the overall consumption response in the credit panel data. In Figure 20, I show the impulse response of vehicle purchases in the credit panel data to a one standard deviation expansionary monetary policy shock in both the BSO series (left) and BRW series (right). The BSO shocks generate the appropriately signed response, while the BRW shocks generate a counter-intuitive consumption response in my sample. In response to expansionary monetary policy shocks, I observe a decline in overall consumption.

Figure 20: Panel Data Consumption Response

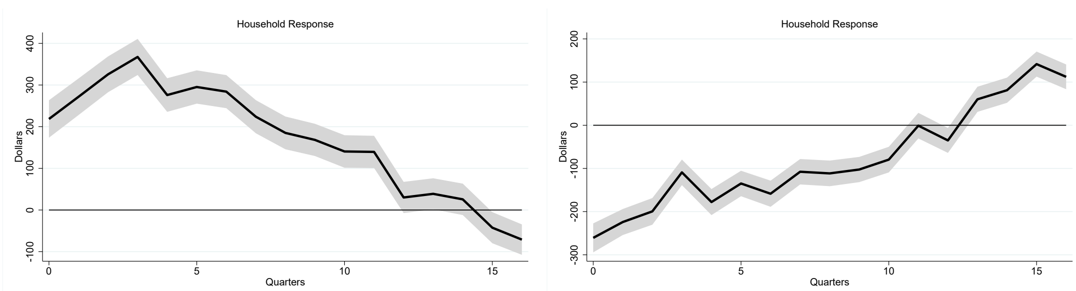

In Figure 21, I also show the impulse responses of various aggregate time series to a one standard deviation expansionary monetary policy shock in both the BSO series (left) and BRW series (right). I include the Federal Funds Rate, GDP, and PCE consumption. The responses of the aggregate time series data are consistent with the results in Figure 20, with the BSO shocks generating a fall in the federal funds rate and a subsequent increase in both consumption and GDP. The BRW shocks do not produce strong consumption or GDP responses over my sample period. Given the responses of the aggregate data, I focus on the Bauer and Swanson (2023) monetary policy shocks for my cross-sectional analysis of the differential monetary policy responses by debt type.

Interest Rate Responses

The Bauer and Swanson shocks use the first four quarterly Eurodollar futures contracts, ED1–ED4, to capture monetary policy surprises, which they rescale so that a one-unit change in the series corresponds to a one percentage point change in the ED4 rate, arguing that Eurodollar futures are a good predictor

Figure 21: Time Series Results: 2004-2020 (Panel Data Sample)

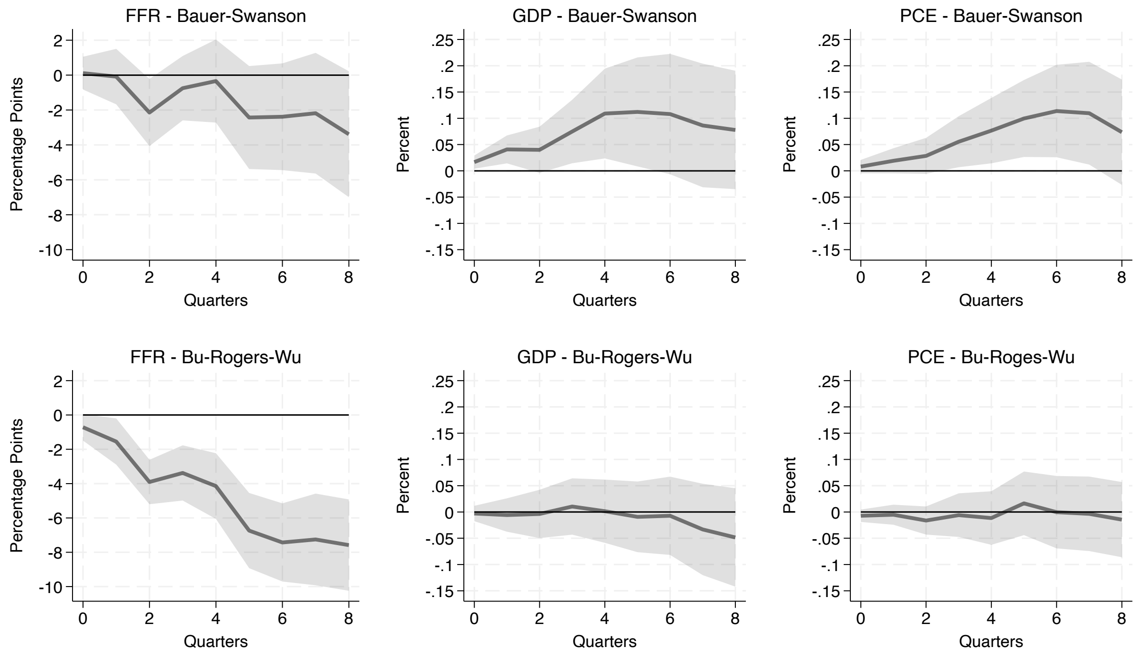

of future values of the federal funds rate. In Table 4 I include the response of interest rates to the BSO monetary policy surprises to verify that the shocks generate the expected movements in rates over the sample period, 2004-2019. The specification is identical to Equation 29, but here I use daily data on the market yield on U.S. Treasury Securities.

Table 4: Mortgage Rate Responses to Monetary Policy Shock

Interest Rate	FFR	2-Yr Treasury	5-Yr Treasury	10-Yr Treasury
MPS	0.348*	0.404***	0.653***	0.597***
	(.232)	(.182)	(.186)	(.203)

Notes: The table shows the change in daily U.S. treasury yields in response to a 1 standard deviation monetary policy shock. Standard errors in parentheses. Source: Federal Reserve Board, retrieved from FRED, Federal Reserve Bank of St. Louis.
* $p<0.32$, ** $p<0.10$, *** $p<0.05$

G Appendix: Robustness Checks - Lag Length

Figure 22: Consumption Response to an Expansionary Monetary Policy Shock by Debt Group Relative to Non-Debtor Response - 4 Lags

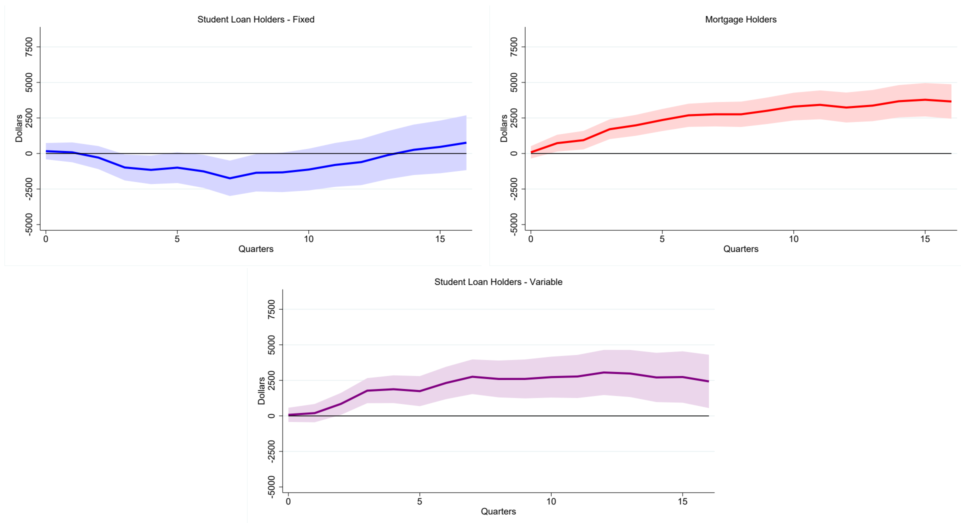

Figure 23: Consumption Response to an Expansionary Monetary Policy Shock by Debt Group Relative to Non-Debtor Response - 8 Lags

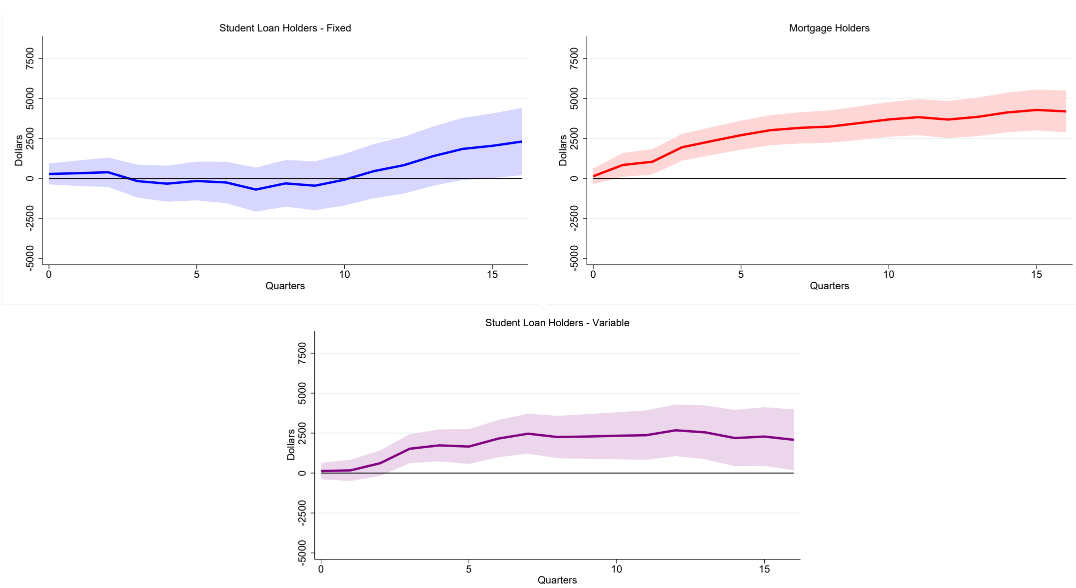

Figure 24: Consumption Response to an Expansionary Monetary Policy Shock by Debt Group Relative to Non-Debtor Response - 12 Lags

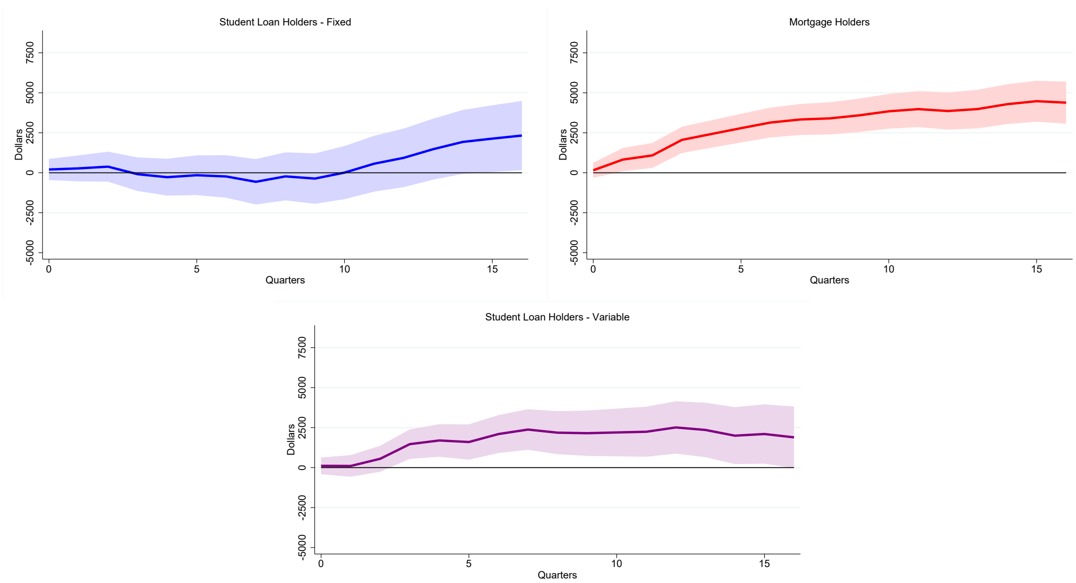

Figure 25: Consumption Response to an Expansionary Monetary Policy Shock by Debt Group Relative to Non-Debtor Response - 16 Lags

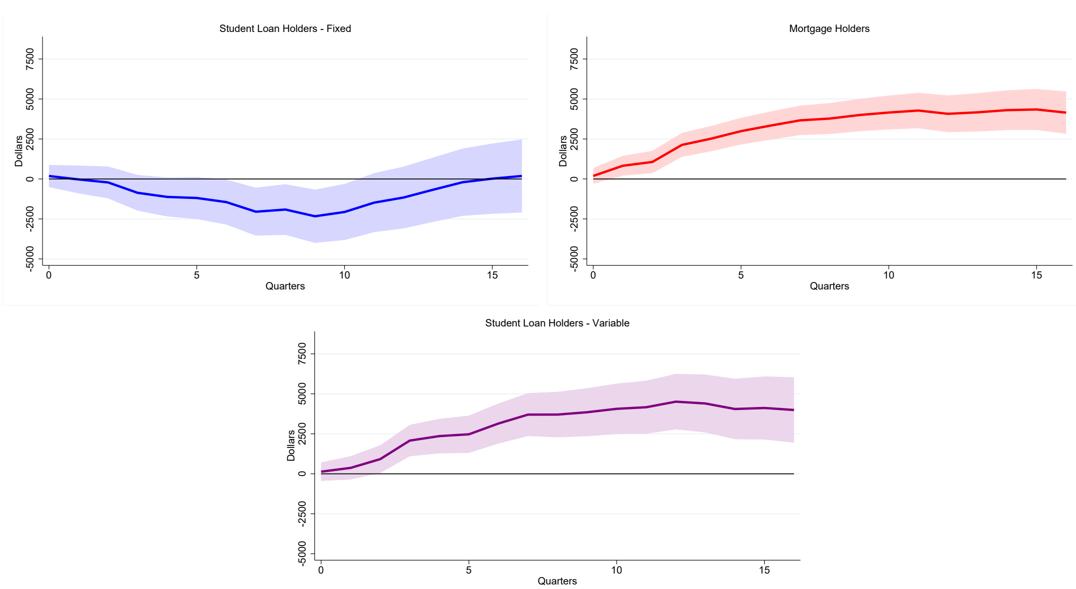

Notes: The figures shows the impulse response of consumption for households with different types of debt relative to the non-debtor response with different lag length controls. Consumption is defined as auto loan originations. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock.

H Appendix: Robustness Checks - Alternate Consumption Measures

One potential concern with the empirical setting is the use of auto loan origination as the primary measure of consumption. As described above, using vehicle purchases as a measure of consumption is not new to this paper. Vehicle purchases capture the dynamics of durable spending well, but in using auto loans to proxy for consumption, I miss out on any changes in expenditures on non-durable goods and services. One way to overcome this lack of non-durable consumption is to use information on credit card balances available in credit reports. For the following results, I re-estimate the main regression in equation 1, but I change the left-hand side variable. Specifically, I utilize the following definition for consumption:

$$\Delta B_{i,t+h} = B_{i,t+h} - B_{i,t+h-1},$$

where $B_{i,t}$ is the total outstanding balance on credit cards in time t.

Using credit card balances in consumer credit data presents three major issues: (1) credit reports capture only a screenshot of the individual's card balances on specific dates potentially missing a large portion of consumption spending, (2) credit bureaus do not collect information on the interest rates on individual credit lines, meaning there is no way to separate accruing interest from current consumption, and (3) measuring consumption using balances may also be affected by the very channel that I am attempting to measure. Mainly, the mechanism for the increased consumption is changes in cash flow deriving from lower interest rates, individuals who receive this increase in cash flow may substitute contemporaneous consumption with paying down old debt, which would show up as a fall in balances for credit cards or other loans.

Credit cards in the US generally serve three purposes to households: consumption smoothing, allocating life-cycle consumption, and means of payment. Fulford and Schuh (2023) find that about 50% of households use credit cards simply as a means of payment and therefore do not carry balances. These households charge their everyday consumption purchases to their credit cards and then pay down the balance over the course of the month. For these households, the main issue with measuring their consumption through credit card balances is that the relatively granular nature of the credit bureau reports means that I might miss a substantial portion of their consumption. Credit card companies generally report account balances to credit bureaus on specific days throughout the month, and these statement days are not always the same as the payment due date for the individual. So if someone chooses to make payments on their credit cards before this or chooses to pay their entire balance, not just the minimum payment or the statement balance, I will miss a significant portion of their consumption purchases.

This issue is enhanced by the quarterly frequency of the data. This would likely cause me to understate the true value of quarterly consumption.

For the other 50% of households who routinely carry a balance and usually pay only the minimum balance, the timing issue is not as severe. However, there is an additional issue with separating out new consumption spending and accrued interest on previous balances. Since credit card companies do not report interest rates or payments, I cannot identify what portion of the changes in credit card balances are coming from interest and what is new consumption.

Figure 26: Cumulative Consumption Response to an Expansionary Monetary Policy Shock by Debt Type Relative to Non-Debtor Response - Consumption Measured using Credit Card Balances

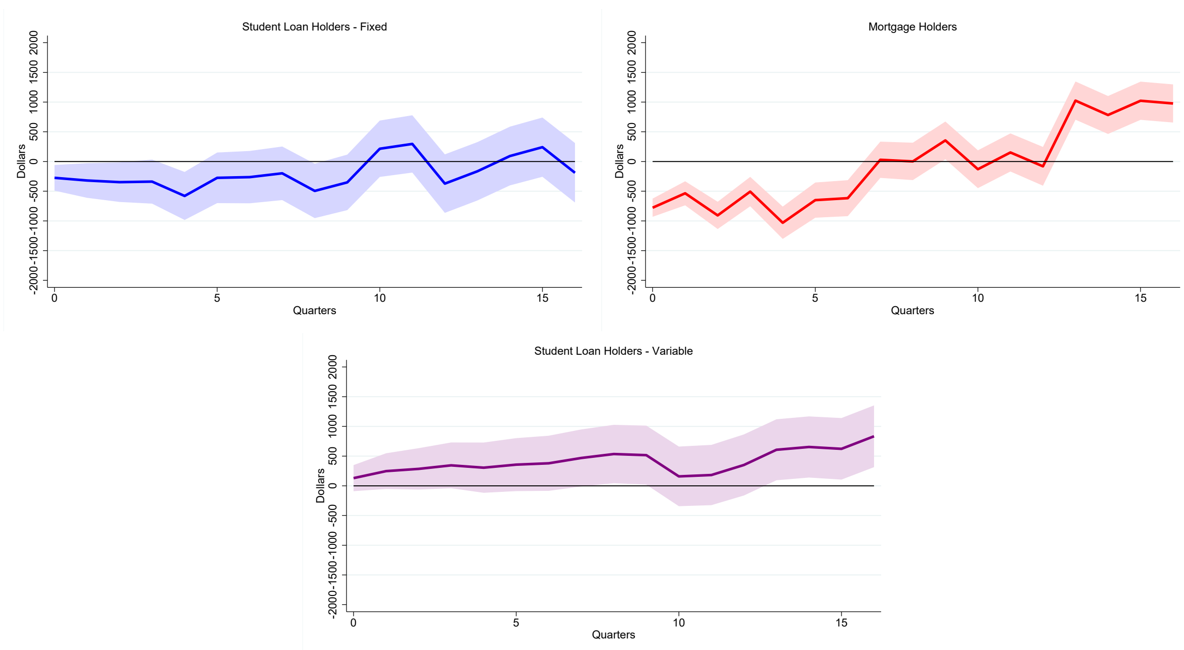

Notes: The figure shows the cumulative impulse response of consumption to an expansionary monetary policy shock for households with different types of debt relative to the non-debtor response. Consumption is measured as the change in credit card balances. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock. Cumulative responses are the integral under the impulse response of the level of consumption.

Figure 26 shows the cumulative interaction effect for the three different types of debt. It appears that both mortgagors and households with fixed-rate student loan debt have slightly lower consumption compared to the non-debtor when measured using credit cards, although the mortgagors eventually increase credit card balances relative to the non-debtor response. Households with variable rate student loan debt have a similar, if not slightly higher, credit card response relative to the non-debtor households.

There is very little that I can do to test the implications of the measurement error discussed above, but I do include some potentially relevant additional results in Figures 28 and 29 below. Including, estimating the same regressions with other loan types, including other secured and unsecured loans,

Figure 27: Cumulative Consumption Response to an Expansionary Monetary Policy Shock by Debt Type Relative to Non-Debtor Response - Consumption Measured Using Changes in All Other Loan Balances

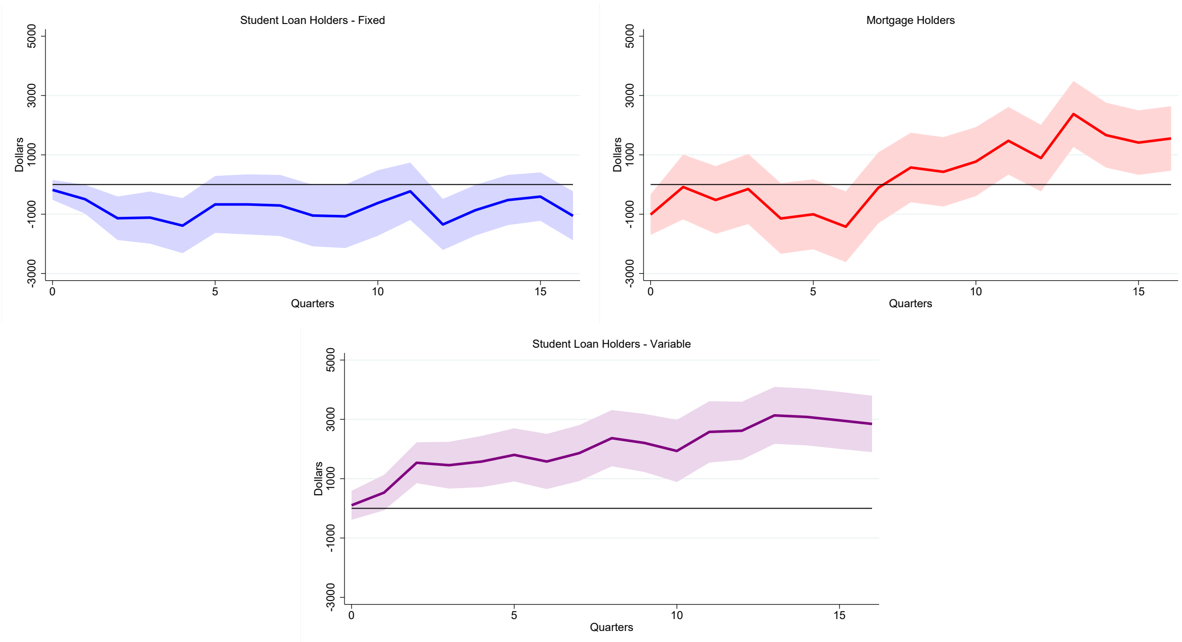

Notes: The figure shows the cumulative impulse response of consumption to an expansionary monetary policy shock for households with different types of debt relative to the non-debtor response. Consumption is measured as the change in all other debt balances, including credit card balances, and other secured and unsecured loan balances. The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock. Cumulative responses are the integral under the impulse response of the level of consumption.

which are not subject to the same measurement issues as credit card balances. Both of these analyses are consistent with the main results that mortgagors and those with variable-rate student loans are more responsive. I also look at the effect on total balances for these other types of debt combined (credit cards, other unsecured, and secured). I estimate the specification with the change in total debt balances, excluding auto loan, mortgage, and student loan debt. I show the results of this specification in Figure 27. Here the results are consistent with the auto loan origination measure. Since the auto loan origination series is better measured and is more likely to capture contemporaneous consumption without interference from other factors that might impact the balance of loans, such as refinancing, changes in payments, etc., I use only the auto loan series in the main specification.

Figure 28: Cumulative Consumption Response to an Expansionary Monetary Policy Shock by Debt Type Relative to Non-Debtor Response - Consumption Measured using Other Secured Loans Balances

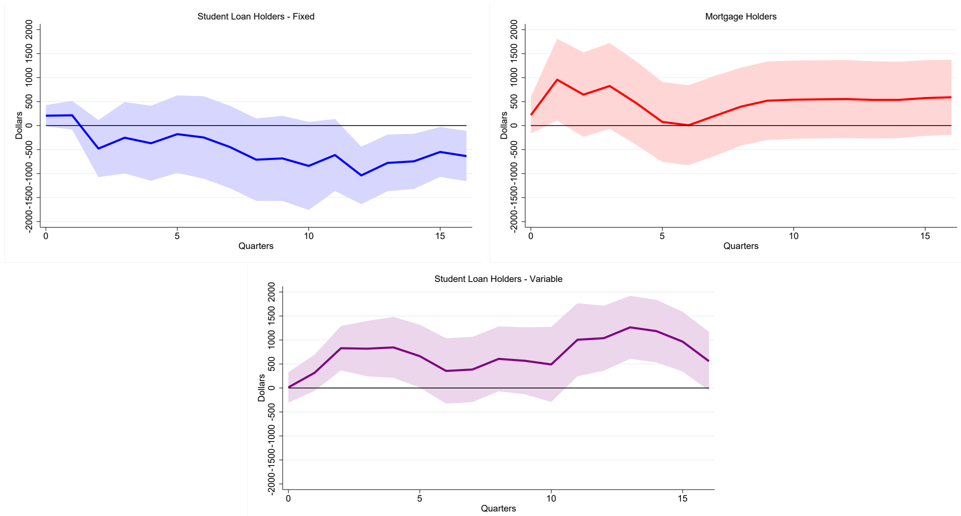

Figure 29: Cumulative Consumption Response to an Expansionary Monetary Policy Shock by Debt Type Relative to Non-Debtor Response - Consumption Measured using Other Unsecured Loans Balances

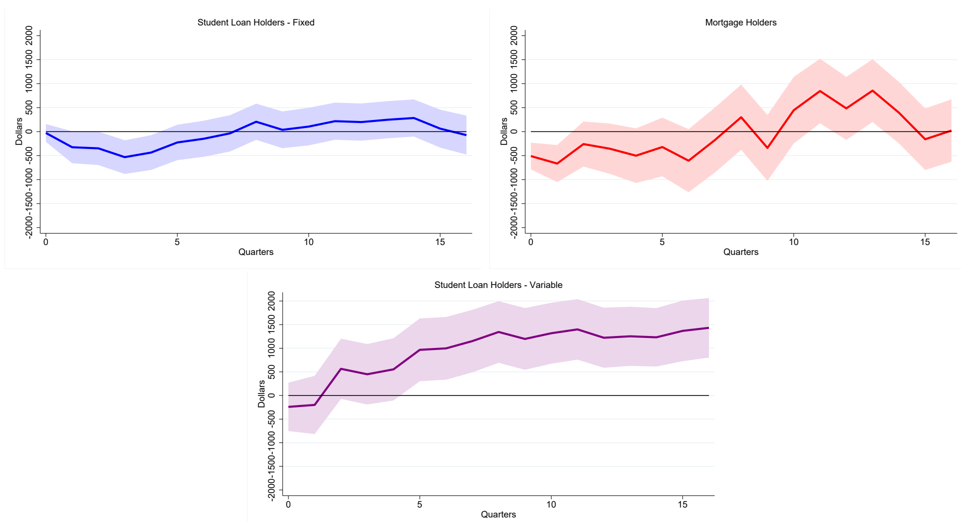

Notes: The figures above show the cumulative impulse response of consumption to an expansionary monetary policy shock for households with different types of debt relative to the non-debtor response. Consumption is measured as the change in other secured loan balances (Figure 28) and other unsecured loan balances (Figure 29). The shock is a one standard deviation decrease in the monetary policy shock series. 90% confidence bands shown. The vertical axis represents 2006 U.S. dollars, and the horizontal axis is quarters following the shock. Cumulative responses are the integral under the impulse response of the level of consumption.

I Appendix: Model Derivations

I.1 General Equilibrium Effect

Equilibrium Conditions

The unknown variables in the model are: C_t, C_t^s, C_t^v, C_t^f, N_t, W_t, P_t, Y_t, C^s. The equations for the individual consumption variables are defined above, and the rest of the conditions are:

$$\begin{aligned}
\text{Production Function: } & Y_t = N_t \text{ for } t \in 1,2 \\
\text{Aggregate Labor Supply: } & \frac{W_2}{P_2} = C_2 N_2^{\chi} \\
\text{Wage Setting: } & W_1 = W_0 \\
\text{Labor Demand: } & \frac{W_t}{P_t} = 1 \text{ for } t \in 1,2 \\
\text{Market Clearing: } & Y_t = C_t \text{ for } t \in 1,2 \\
\text{Total Consumption: } & C_t = C_t^s + C_t^v + C_t^f \text{ for } t \in 1,2 \\
\text{Total Labor: } & N_t = N_t^s + N_t^v + N_t^f = (1-\epsilon)N_t + \epsilon(1-\lambda)N_t + \epsilon\lambda N_t \text{ for } t \in 1,2
\end{aligned}$$

The remaining variables, i_0, i_1, ϵ, λ, and χ are exogenous.

Aggregate Demand

To derive aggregate demand, I start with the definition of aggregate consumption. Using the market clearing conditions $Y_1 = C_1$, and $W_1 = P_1$, I get the following expression for first period output:

$$C_1 = C_1^s + C_1^v + C_1^f$$
$$Y_1 = C_1^s + C_1^v + C_1^f$$

$$Y_1 = \frac{1}{(1+\beta)} \frac{1}{P_1} \left[W_1 N_1^s + i_0 D^v + \bar{i} D^f + \frac{W_2 N_2^s + D^v + D^f}{(1+i_1)} \right] + \frac{1}{P_1} \left[W_1 N_1^v - i_0 D^v \right] + \frac{1}{P_1} \left[W_1 N_1^f - \bar{i} D^f \right] \quad (32)$$

From the firm's first order condition we have $W_1 = P_1$, so the real wage $\frac{W_1}{P_1} = 1$. Now taking the labor market clearing condition $N_1^s = (1-\epsilon)N_1$, $N_1^v = \epsilon(1-\lambda)N_1$, and $N_1^f = \epsilon\lambda N_1$ and the production function, $Y_1 = N_1$,

$$Y_1 = \frac{1}{(1+\beta)} \left[(1-\epsilon)Y_1 + i_0 \frac{D^v}{P_1} + \bar{i}\frac{D^f}{P_1} + \frac{1}{(1+i_1)} \left[\frac{W_2 P_2}{P_2 P_1}(1-\epsilon)Y_2 + \frac{D^v}{P_1} + \frac{D^f}{P_1} \right] + \epsilon(1-\lambda)Y_1 - i_0\frac{D^v}{P_1} + \epsilon\lambda Y_1 - \bar{i}\frac{D^f}{P_1} \right.$$

$$Y_1 - \frac{(1-\epsilon)}{(1+\beta)} Y_1 - \epsilon Y_1 = \frac{1}{(1+\beta)} \left[i_0 \frac{D^v}{P_1} + \bar{i}\frac{D^f}{P_1} + \frac{1}{(1+i_1)} \left[w_2 \Pi_2 (1-\epsilon) Y_2 + \frac{D^v}{P_1} + \frac{D^f}{P_1} \right] - i_0 \frac{D^v}{P_1} + \bar{i}\frac{D^f}{P_1} \right. .$$

Now I have an expression for Y_1 in terms of exogenous parameters, the interest rates, and second period income.

$$Y_1 = \frac{1}{(1-\frac{1}{1+\beta})(1-\epsilon)}\left[\frac{1}{1+\beta}\left[\frac{i_0 D^v}{P_1} + \frac{\bar{i}D^f}{P_1} + \frac{1}{(1+i_1)}\left[w_2\Pi_2(1-\epsilon)Y_2 + \frac{D^v}{P_1} + \frac{D^f}{P_1}\right]\right] - i_0\frac{D^v}{P_1} + \bar{i}\frac{D^f}{P_1}\right] \quad (33)$$

In order to pin down second period output and income, I assume flexible prices and wages and solve for the natural level of output.

From the firm's first-order conditions,

$$\frac{W_2}{P_2} = w_2 = 1 \quad (34)$$

firms hire labor until the marginal product is equal to the real wage.

Combining Equation 34 and the aggregate labor supply curve (derived below) $\frac{W_2}{P_2} = C_t N_t^{\chi}$ I get,

$$C_2 N_2^{\chi} = 1$$

Then combining this with the market clearing condition $Y_2 = C_2$ and the production function $Y_2 = N_2$ gives, $N_2 = 1$. Thus, output in period two will be equal to

$$Y_2 = 1.$$

Finally, to pin down second period inflation, I assume the central bank follows a simple monetary policy rule implementing an inflation target, $P_2 = \Pi P_1$. Therefore, second period output and inflation are independent of the shock to the monetary policy in period 0.

I.2 Aggregate Labor Supply

The individual labor supply curves for the three agents are:

$$\theta^s C_2^s (N_2^s)^{\chi} = \frac{W_2}{P_2}$$

$$\theta^v C_2^v (N_2^v)^{\chi} = \frac{W_2}{P_2}$$

$$\theta^f C_2^f (N_2^f)^{\chi} = \frac{W_2}{P_2}$$

Assuming $N_s = (1-\epsilon)N_t$, I can set $\theta^s = \frac{n}{(1-\epsilon)^\chi}$, where $n = 3$, the total number of agents in the economy. The labor supply curve simplifies to,

$$\frac{3}{(1-\epsilon)^\chi}(1-\epsilon)^\chi N_2^\chi C_2^s = 3N_2^\chi C_2^s = \frac{W_2}{P_2} \tag{35}$$

Similarly, using $N_v = \epsilon(1-\gamma)N_t$, and $N_v = \epsilon\gamma N_t$ and setting θ^v and θ^f accordingly, results in the following:

$$\frac{3}{(\epsilon(1-\gamma))^\chi}(\epsilon(1-\gamma))^\chi N_2^\chi C_2^v = 3N_2^\chi C_2^v = \frac{W_2}{P_2} \tag{36}$$

$$\frac{3}{(\epsilon\gamma)^\chi}(\epsilon\gamma)^\chi N_2^\chi C_2^f = 3N_2^\chi C_2^f = \frac{W_2}{P_2} \tag{37}$$

The aggregate labor supply curve is simply the sum of the individual labor supply curves.

$$3N_2^\chi C_2^s + 3N_2^\chi C_2^v + 3N_2^\chi C_2^f = 3(\frac{W_2}{P_2})$$

$$3N_2^\chi (C_2^s + C_2^v + C_2^f) = 3(\frac{W_2}{P_2})$$

$$N_2^\chi C_2 = \frac{W_2}{P_2} \tag{38}$$