The Transmission of Monetary Policy under the Microscope

Martin Blomhoff Holm, Pascal Paul, Andreas Tischbirek*

Abstract

We investigate the transmission of monetary policy to household consumption using administrative data on the universe of households in Norway. Based on identified monetary policy shocks, we estimate the dynamic responses of consumption, income, and saving along the liquid asset distribution of households. For low-liquidity but also for high-liquidity households, changes in disposable income are associated with a sizable consumption reaction. The impact consumption response is closely linked to interest rate exposure, which is negative at the bottom but positive at the top of the distribution. Indirect effects of monetary policy gradually build up and eventually outweigh the direct effects.

Keywords: Monetary policy, Administrative data, Household balance sheets, Liquidity constraints, Heterogeneous agent New Keynesian models

JEL Codes: D31, E21, E43, E52, G51

*We thank Adrien Auclert, Florin Bilbiie, Corina Boar, James Cloyne, Andreas Fagereng, Kilian Huber, Òscar Jordà, Greg Kaplan, Ralph Luetticke, Ben Moll, Chris Moser, Emi Nakamura, Christina Patterson, Matthew Rognlie, Morten Ravn, Jón Steinsson, and seminar participants at UC Berkeley, the University of Chicago, Columbia Business School, the University of Bonn, UCL, HU Berlin, the Federal Reserve Board, Norges Bank, the NBER Summer Institute, the "Monetary Policy and Heterogeneity" conference at the Federal Reserve, the EABCN conference on "Empirical Advances in Monetary Policy", and the World Congress of the Econometric Society for their insights. We also thank Anne Sofie Jore for helping us with the collection of archive data from Norges Bank and Chitra Marti, Ben Shapiro, Dori Wilson, and Winnie Yee for excellent research assistance. All errors are our own. This article has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 851891). We also thankfully acknowledge funding from the Swiss National Science Foundation. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. Holm: University of Oslo, email: martin.b.holm@outlook.com. Paul: Federal Reserve Bank of San Francisco, email: pascal.paul@sf.frb.org. Tischbirek: HEC Lausanne, University of Lausanne, email: andreas.tischbirek@unil.ch.

1 Introduction

What is the effect of monetary policy on consumption, income, and saving across the distribution of households? And what role does household heterogeneity play in the transmission mechanism? A new class of heterogeneous agent models with nominal rigidities has recently emerged that provides valuable insights into these questions.[1] However, direct empirical evidence on the responses to monetary policy at the micro level is scarce. In this paper, we draw on detailed administrative data to investigate the transmission of monetary policy to household consumption and examine the empirical validity of the key predictions implied by theory.

Recent Heterogeneous Agent New Keynesian (HANK) models have several appealing features in comparison with standard Representative Agent New Keynesian (RANK) models. In RANK frameworks, consumption-saving behavior is generally closely in line with the permanent income hypothesis (Kaplan and Violante, 2014; Bilbiie, 2020). This implies that the marginal propensity to consume (MPC) associated with temporary income changes is very small, a feature that is at odds with empirical estimates.[2] Almost the entire consumption response to monetary policy is therefore due to direct, partial equilibrium, effects that operate largely through intertemporal substitution. In contrast, HANK models are capable of generating MPCs that are closer to the values estimated from empirical data. In the presence of borrowing costs or constraints, income changes of households with low wealth are imperfectly smoothed and permitted to feed into consumption. Kaplan and Violante (2014) show that this continues to be the case for households with large illiquid but small liquid wealth if holdings of illiquid assets are subject to adjustment costs.[3] In a two-asset model that can give rise to a realistic fraction of households with few liquid assets and therefore to a sizable MPC at the aggregate level, Kaplan, Moll and Violante (2018) demonstrate that interest rate changes influence household consumption predominantly through indirect, general equilibrium, effects by affecting the disposable income of households.

Guided by HANK models, we provide a detailed empirical account of the effects of monetary policy at the household level. In line with theory, our focus lies on the role played by household balance sheets, in particular liquid asset positions, and the relative importance of direct and indirect effects. Empirical evidence for the monetary transmission mechanism implied by HANK models has proved difficult to obtain. To investigate the micro-level responses to policy rate changes, a panel data set is required that spans many years and includes detailed information on the balance

[1] Several contributions extend the standard Bewley-Huggett-Aiyagari incomplete markets framework to include nominal rigidities. Examples are Oh and Reis (2012), McKay and Reis (2016), Den Haan, Rendahl and Riegler (2017), Guerrieri and Lorenzoni (2017), Ravn and Sterk (2017) and Bayer et al. (2019), among many others. Detailed analyses of monetary policy are contained in Gornemann, Kuester and Nakajima (2016), McKay, Nakamura and Steinsson (2016), Kaplan, Moll and Violante (2018), Auclert (2019), Hagedorn et al. (2019), Challe (2020), Kekre and Lenel (2020), and Luetticke (2020), for example. Influential tractable models are developed in Campbell and Mankiw (1989), Mankiw (2000), Galí, López-Salido and Vallés (2007), Bilbiie (2008), Debortoli and Gali (2018), and Bilbiie (2020).

[2] See, e.g., Johnson, Parker and Souleles (2006), Jappelli and Pistaferri (2010), Parker, Souleles, Johnson and McClelland (2013), Misra and Surico (2014), Bunn, Roux, Reinold and Surico (2018), and Fagereng, Holm and Natvik (2021).

[3] Kaplan, Violante and Weidner (2014) provide empirical support for the existence of "wealthy hand-to-mouth households" in a number of countries.

sheets, income, and consumption of households. We therefore turn to a country that collects rich information from its inhabitants: Norway. Specifically, we draw on administrative data that contain records of the income and wealth of all households registered in Norway between 1996 and 2015. Using the information contained in the administrative data, we impute the consumption expenditures of households based on their budget identity. Equipped with this data set, we are able to give a comprehensive picture of the consumption response to monetary policy shocks and its determinants at fine segments along the liquid asset distribution.

To overcome monetary policy endogeneity, we derive a series of identified monetary policy shocks for Norway using the approach in Romer and Romer (2004). Before turning to the micro data, we estimate the responses to these shocks using aggregate data. We obtain textbook impulse responses across a wide range of macroeconomic aggregates: indicators for economic activity contract and prices fall after a monetary tightening. For the time series, which are available at different frequencies, the impulse responses show similar patterns independent of whether monthly, quarterly, or annual data are used in the estimation. The shape and stability of the responses across the different variables and frequencies give us confidence in the identification and in the analysis at the household level, for which we confront the shocks with the administrative data that are available at an annual frequency.

We then turn to analyzing the micro-origins of the macro responses. In congruence with the HANK literature, we divide the population of households into groups of equal size according to their location in the liquid asset distribution and estimate a separate set of impulse responses for each segment of the distribution. The consumption responses move closely with the income responses across the entire distribution. Comparing households at the bottom of the distribution with households around the median reveals differences in the consumption-saving behavior though. When disposable income begins to fall in response to a monetary policy contraction, households with low liquid asset holdings let their consumption decline, while households with intermediate amounts of liquid assets initially reduce saving or increase borrowing, as predicted by theory.

For households at the top of the liquid asset distribution, our results are not consistent with the predictions of standard HANK models. We find that households with large liquidity positions increase consumption substantially in response to a monetary tightening and that the increase on impact is related to a rise in interest income, which is directly affected by the policy rate. While HANK models can generate large and positive consumption responses to an interest rate increase for liquidity-rich households if the income effect dominates the substitution effect, they generally do not produce the large consumption response *relative* to the income response that we observe. In contrast to these models, our estimates indicate that households at the top of the liquid asset distribution have sizable MPCs out of the temporary changes in disposable income induced by monetary policy.

Interest income and expenses play a significant role in shaping the response of disposable income to monetary policy shocks. The effect is particularly strong for households at both ends of the liquid asset distribution. To isolate the cash-flow channel associated with interest-sensitive asset and debt positions, we reorder households according to their net interest rate exposure, a measure closely related to the concept of "unhedged interest rate exposure" by Auclert (2019). The net financial income response of households at the bottom (net borrowers) is the mirror image of that of households at the top (net creditors) and income changes are permitted to affect household consumption at both ends of the distribution. Several factors contribute to the pronounced cash-flow effects that we observe. The ten percent of households with the highest net interest rate exposure hold large amounts of deposits with a group median across all years of about 100,000 U.S. dollars, more than twice the median annual household income after taxes in Norway. Households at the bottom of the distribution tend to have substantial mortgage debt. Since deposit rates are comparably elastic to the policy rate and around 90 percent of mortgage contracts carry an adjustable rate, interest income and expenses are highly responsive to monetary policy.

Our estimates also uncover strong ties between the responses of consumption and nonfinancial income, suggesting that general equilibrium effects play an important role for the transmission of monetary policy as argued by Kaplan, Moll and Violante (2018). We exploit the data to provide empirical estimates of the relative importance of direct and indirect effects. The relative size of these effects depends on the impulse response horizon. On impact, the aggregate consumption response is almost entirely driven by direct effects. Around two years after the shock, at the time when the nonfinancial income response builds up, the indirect effects start to dominate the direct effects. Quantitatively, the importance of the indirect effects is of comparable size as in the HANK model by Kaplan, Moll and Violante (2018), with the difference that they unfold only several years after the shock.

A growing empirical literature inspects the transmission mechanism of monetary policy to household consumption. Cloyne, Ferreira and Surico (2020) use survey data to investigate whether differences in the balance sheets of households, approximated by their housing tenure status, affect the consumption response to monetary policy shocks in the United States and the United Kingdom. They find that households with a mortgage respond more strongly than outright home owners and renters. The findings are interpreted as evidence that liquidity positions play a key role for the consumption response of households, since mortgagors tend to have low liquid asset holdings. A benefit of our data is that they allow us to observe liquidity positions directly and therefore to separate households along this dimension.

A number of contributions emphasize the importance of mortgage contracts for the pass-through of monetary policy to household consumption. In the United States, the majority of mortgages carries a fixed rate. Households can reduce interest expenses or extract housing equity if they refinance their loans in response to rate cuts. Wong (2019) finds that the consumption response to monetary policy is stronger among households that adjust their mortgages than among those that do not. Di Maggio et al. (2017), exploiting specifics of the mortgage design in the United States, and Floden et al. (2020), using administrative data from Sweden, reach similar conclusions about the effects exerted by adjustable-rate mortgages. The propensity to refinance has been linked to household age and loan size (Wong, 2019), the path of past interest rates (Berger et al., 2018; Eichenbaum, Rebelo and Wong, 2019), and housing equity (Beraja, Fuster, Hurst and Vavra, 2018). Our results for households with negative net interest rate exposure confirm that the cash flows associated with mortgages play an important role in monetary policy transmission. In addition, we highlight the significance of net creditors and estimate dynamic responses for several years after a shock. The latter is of considerable interest, since consumption has been shown to respond to monetary policy shocks with a sizable delay (see, e.g., Christiano, Eichenbaum and Evans, 2005). Our emphasis on the dynamic effects of policy changes in the context of household heterogeneity is shared by Auclert, Rognlie and Straub (2018), although their focus lies on fiscal policy.

Several additional papers estimate key moments in the data and use those moments as model inputs to study the channels through which monetary policy affects household consumption. Auclert (2019) decomposes the aggregate consumption response to monetary policy into different channels to evaluate the role played by redistribution in the presence of heterogeneity in MPCs. Based on survey data, he concludes that redistribution amplifies the aggregate response. Crawley and Kuchler (2020) refine these estimates using administrative data from Denmark. Slacalek, Tristani and Violante (2020) build on Auclert's decomposition to separate direct from indirect effects. Patterson (2018) relies on a similar method to study the amplification of shocks and the severity of recessions. Auclert, Rognlie and Straub (2020) also decompose the total response of monetary policy into direct and indirect effects. They find that in their setup with inattentive households the indirect effects far outweigh the direct effects.

The remainder of the paper is organized as follows. The next section derives the series of monetary policy shocks and computes aggregate responses to these shocks at different frequencies. Section 3 describes the data, discusses the consumption imputation, and presents descriptive statistics. Section 4 contains our main results on the transmission of monetary policy at the household level. Section 5 decomposes the aggregate consumption response into direct and indirect effects. Section 6 concludes and discusses implications of our findings for HANK models.

2 Monetary Policy Identification

Most of the variation in central bank interest rates is due to the systematic response of policy to current or expected future economic conditions. To identify the causal effects of monetary policy, it is therefore necessary to isolate shifts in monetary policy instruments that are orthogonal to policy responses to the behavior of the economy. In this paper, we rely on the approach by Romer and Romer (2004) to identify monetary policy shocks.[4] The idea of this approach is to orthogonal-

[4] A popular alternative is to use financial markets data to extract surprise changes in interest rates around policy announcements (see, e.g., Kuttner, 2001, Guerkaynak, Sack and Swanson, 2005, and Gertler and Karadi, 2015). However,

ize policy rate changes against the central bank's forecasts of its macroeconomic targets in a first step. The estimated residuals serve as a measure of monetary policy shocks. In a second step, the externally identified shock series can be employed to estimate impulse responses.

The key policy rate of Norges Bank, the Norwegian central bank, is the sight deposit rate. Its historical evolution is shown in Figure A.1 in Appendix A.1. As can be seen from the figure, the policy rate never touched the zero-lower bound over the entire interval considered. We are therefore able to study the effects of conventional policy rate changes in recent years without having to account for a period in which the policy rate was constrained. On a policy-meeting frequency, we estimate

$$\Delta i_m = \alpha_1 + \alpha_2 i_{m,-1} + \sum_{k=0}^{1} \beta_k^{\pi} \pi_{m,t+k} + \sum_{k=0}^{1} \beta_k^{\Delta\pi} \Delta\pi_{m,t+k} + \sum_{k=0}^{1} \beta_k^{y} y_{m,t+k} + \sum_{k=0}^{1} \beta_k^{\Delta y} \Delta y_{m,t+k} + \gamma_1 ex_{m,-1} + \gamma_2 I_m^{IT} \cdot ex_{m,-1} + \epsilon_m^{MP} \quad , \tag{1}$$

where Δi_m is the change of the policy rate at meeting m and $i_{m,-1}$ is the level of the policy rate prior to meeting m. Meeting m takes place in period t.[5] Following Romer and Romer (2004), we include central bank forecasts for GDP $y_{m,t+k}$ and the CPI $\pi_{m,t+k}$ for horizon $t+k$ and the corresponding forecast changes, denoted $\Delta\pi_{m,t+k}$ and $\Delta y_{m,t+k}$. We rely on historical forecasts from Norges Bank for all policy meetings, whenever these were constructed shortly before a meeting. When this is not the case, we follow Cloyne and Huertgen (2016) in using forecasts by market participants to proxy for the forecasts of the central bank. Appendix A.2 describes these forecasts further and gives a detailed protocol for their assignment to the policy meetings.

The specification in (1) deviates from the one employed by Romer and Romer (2004) in three ways. First, we use annual forecasts for the current and the next year as opposed to quarterly ones since these are available for a relatively long historical sample.[6] Second, we do not include a contemporaneous forecast for the unemployment rate since such a forecast is not available for a longer historical sample. Third, we also account for a switch in policy regimes over our sample. From March 2001 onward, Norges Bank officially committed itself to an inflation targeting regime. In the years before, the central bank also targeted the exchange rate. We therefore include as additional explanatory variables the level of the exchange rate on the day before the meeting $ex_{m,-1}$ and the same variable interacted with an indicator variable I_m^{IT} that takes the value of one for the pre-inflation targeting era.[7] The residual ϵ_m^{MP} in regression (1) is a measure of the monetary policy shock associated with meeting m.

for Norway, it is difficult to reconstruct with certainty at which points in time the information about policy decisions was transmitted to financial markets in the early years of the sample.

[5] Occasionally, there are multiple policy rate changes shortly after one another. We combine policy rate changes within one month and apply the date of the later meeting to the combined rate change. We checked that there are no such instances across months. The results are nearly identical without this adjustment.

[6] The earliest quarterly forecasts for the output gap and a price index start in late 2005, which would restrict the analysis to the second half of the sample.

[7] We use historical data of the import-weighted exchange rate from Norges Bank.

2.1 Estimation Results

We estimate (1) using ordinary least squares for the sample 1994:M1-2018:M12. The results are shown in Table 1. The estimated coefficients have the expected signs and the ones associated with the forecasts and the forecast changes that are statistically different from zero are all positive. That is, if projected inflation or output growth is high or has been increasing relative to the prior comparison forecast, monetary policy tightens to lean against these macroeconomic developments. Moreover, the constant and the coefficient on the lagged policy rate are negative, reflecting the secular decline in interest rates over our sample and a mean-reversion in policy rates, respectively. However, only the coefficient on the lagged policy rate is statistically different from zero. The estimates and R^2 of around 0.3 are consistent with the findings of Romer and Romer (2004) and Cloyne and Huertgen (2016). The estimated coefficients associated with the exchange rate turn out to be not statistically significant, even though their positive signs imply that monetary policy tightens if the currency is weak before the meeting.

Table 1: **Determinants of Changes in the Policy Target Rate.**

	Constant	$i_{m,-1}$	π_m	y_m	$ex_{m,-1}$	$I_m^{IT} \cdot ex_{m,-1}$
	-0.50 (0.22)	-0.02* (0.09)			0.02 (0.95)	0.06 (0.34)
Current Year			0.06** (0.04)	0.05 (0.37)		
Next Year			0.04 (0.44)	0.04 (0.62)		
ΔCurrent Year			0.02 (0.28)	0.27*** (0.00)		
ΔNext Year			0.11** (0.02)	-0.04 (0.58)		
$N = 162$						
$R^2 = 0.30$						

Notes: Estimation results for (1). Sample: 1994:M1-2018:M12 (excluded: 1998:M8, 2008:M10/M12, see Appendix A.2 for details), p-values in parentheses, *** p<0.01, ** p<0.05, * p<0.1.

Similar to Romer and Romer (2004), we convert the series of residuals ϵ_m^{MP} in (1) from a meeting frequency into a monthly, quarterly, and annual time series ϵ_t^{MP} by assigning each shock to the month, quarter, or year in which it occurred. If there are multiple meetings within a period, then we aggregate the associated shocks by summing up the shocks within that time period. If there are no policy meetings, the corresponding shock is set to zero. The monthly series of monetary policy shocks is shown in Figure 1, while the quarterly and annual shocks are shown in Figures A.2 and A.3 in Appendix A.3.

Several properties are worth mentioning. First, some shocks are large: a few are more than 50 basis points, reflecting the overall large movements of the policy rate over our sample period (see Figure A.1 in Appendix A.1). Romer and Romer (2004) find similar large movements in the shock

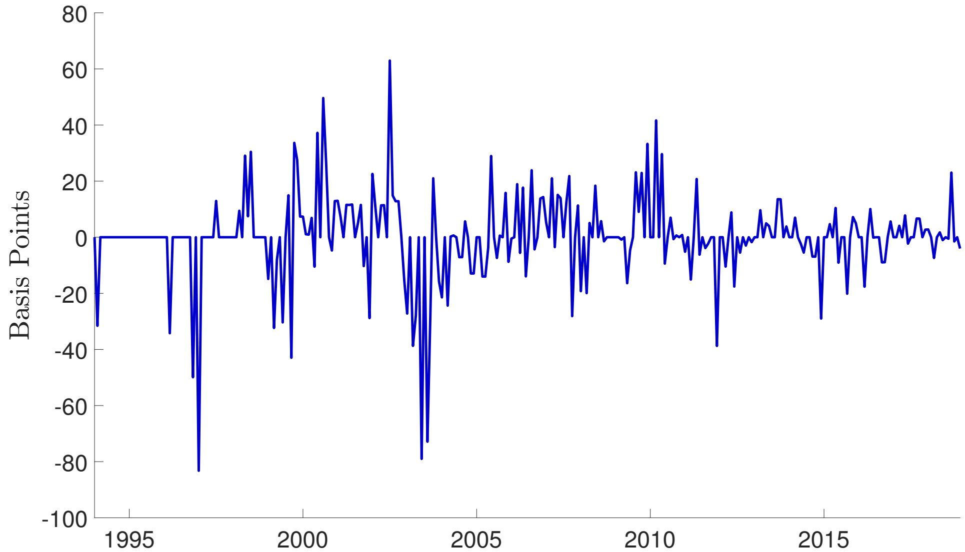

Figure 1: **Monthly Series of Monetary Policy Shocks.**

series for the United States. Particularly noticeable are the contractionary shocks around 2002 followed by the sharp easing shocks in 2003. The related policy rate changes were criticized by external observers who detected "policy mistakes that kept monetary policy too tight at the end of 2002 and early in 2003 and perhaps too loose at the end of 2003" (Bjørnland et al., 2004). The identification approach picks up both of these movements as relatively unrelated to the typical response to contemporaneous forecasts. In Appendix A.4, we add more texture to the identification and provide a quantitative analysis of the largest shocks that we uncover. Second, the shocks become smaller toward the end of the sample, in line with the reduced volatility of the policy rate in recent years. Third, in Figure A.4 in Appendix A.3, we compare the shock series with the actual rate changes. While the two tend to move in the same direction, as in Romer and Romer (2004), there are often significant differences. Fourth, we test whether the monthly shock series is predictable based on past data. Similar to Coibion (2012) and Cloyne and Huertgen (2016), we use lagged monthly changes of the unemployment rate, percentage changes in industrial production, and consumer price inflation as predictors. We find no evidence of predictability (see Table A.2 in Appendix A.5).

2.2 Impulse Responses - Macro Aggregates

Based on the identified shocks, we run a series of local projections on a monthly, quarterly, and annual frequency. Let y_t be the outcome variable at time t, e.g., (log) real GDP or the unemployment rate. Following Jordà (2005), we estimate

$$y_{t+h} - y_{t-1} = \alpha^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{k=1}^{K} \gamma_k^h X_{t-k} + u_t^h \quad , \tag{2}$$

where $h = 0, 1, .., 5$ for annual data, $h = 0, 1, .., 20$ for quarterly data, and $h = 0, 1, .., 60$ for monthly data. The estimated coefficients β^h give the percentage (point) change at horizon h in response to a 100-basis-point monetary policy shock at the respective frequency.[8] Note that we leave the contemporaneous response unrestricted, in contrast to a typical Cholesky identification. X_t denotes a vector of controls. Our specification includes three years of lagged values of the monetary policy shock as in Romer and Romer (2004), but we do not add lagged values of the dependent variable or any other variable as regressors.[9]

We consider a wide range of outcome variables at the aggregate level. Table A.3 in Appendix A.6 gives precise details on the time series employed.[10,11] The estimated impulse responses to a contractionary shock of 100 basis points are shown in Figure 2 (monthly frequency), Figure 3 (annual), and Figure A.8 (quarterly) in Appendix A.7. Across the different frequencies, we obtain textbook responses to a monetary tightening. The policy rate increases and subsequently reverses, a feature that we return to below when we analyze the responses based on the micro data. Economic activity contracts as the unemployment rate rises and industrial production, GDP, and consumption expenditures fall. Consumer prices and real wages and salaries decline.[12] Figures A.7-A.11 in Appendix A.7 show the responses of a number of additional variables. Throughout, the responses have the expected signs: production and investment measures decline; various price indices, including house and stock prices, fall; hours worked decline; household income falls; and measures of income inequality increase, consistent with the findings by Coibion et al. (2017) for the United States. By and large, the estimated responses are statistically different from zero at the 95 or 68 percent confidence level. Further, in Appendix A.8, we check and confirm the robustness of our results to various modifications of the baseline regressions and compare the impulse responses to the analogous ones based on U.S. data, which turn out to be very similar.

Importantly, the impulse responses share a similar dynamic shape across the different frequencies. For example, the unemployment rate rises steadily in response to a monetary tightening, with a peak response after around 3.5 years, and falls thereafter. However, the size of the re-

[8]Throughout, we interpret the estimated shocks as direct observations of the structural monetary policy shocks as opposed to instruments (see, e.g., Mertens and Ravn, 2013; Stock and Watson, 2018). If our estimated shocks are imperfectly correlated with the true structural shocks, then the local projections in (2) are still valid. However, the impulse responses should then be interpreted as relative impulse responses, as opposed to absolute ones (see Paul, 2020, for details).

[9]To choose the lag length for the monetary policy shocks, we consult the Akaike and Bayesian information criteria. Across various outcome variables, the information criteria tend to favor longer than three lags of the monetary policy shock for near-term impulse responses and shorter ones for impulse responses further out. The chosen lag length is therefore a reasonable unifying compromise across variables and impulse response horizons. We further test (and confirm) the robustness of the results to the chosen lag length in Appendix A.8. We do not include lagged shocks as controls in the equation for the policy rate, since the policy rate responds on impact to the shock. The responses of the policy rate are largely unaffected by including additional controls.

[10]Most series are obtained from Statistics Norway. They are generally denominated in real units, seasonally adjusted, and provided for the full length of the sample. If not, we adjust the series using consumer prices, add seasonal dummies as additional regressors to equation (2), or estimate the local projections for the longest possible sample.

[11]The industrial production series shows a structural break over the years considered, which is due to a change in oil extraction, as illustrated in Figure A.6 in Appendix A.6. To account for this break, we consult a Chow test and, based on the result, include an additional dummy that equals one pre-2002:M2 and zero otherwise into the respective regressions.

[12]We use the CPI-AEL as our main indicator for consumer prices. This index excludes electricity and thereby mitigates the influence of the energy production sector (including oil and gas) on prices. We find similar responses using the overall CPI and several other consumer price indices supplied by Statistics Norway.

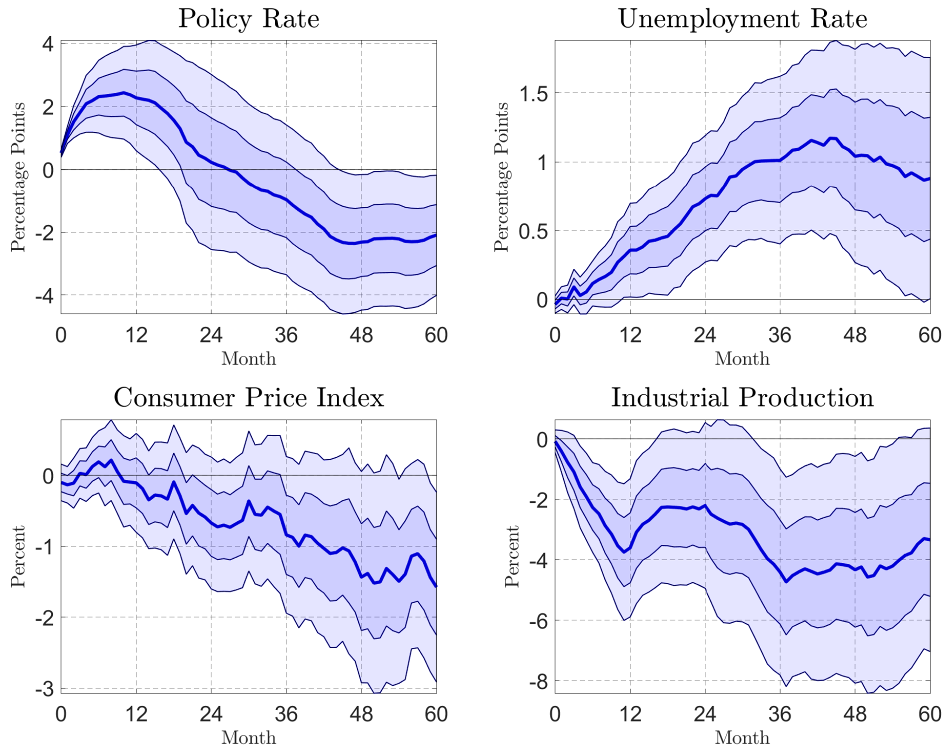

Figure 2: **Impulse Responses at a Monthly Frequency.**
Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at a monthly frequency, based on the local projection approach in (2). 95 and 68 percent confidence bands shown, using Newey and West (1987) standard errors. Table A.3 in Appendix A.6 lists the data sources. Additional impulse responses at a monthly frequency are shown in Figure A.7 in Appendix A.7.

sponses differs somewhat across frequencies. The response of the unemployment rate peaks at around 1.2 percent at the monthly frequency, 1.0 percent at the quarterly frequency, and 0.7 percent at the annual frequency. The policy rate equally increases by less at the annual frequency. In Appendix A.9, we show that the responses across different frequencies are of similar size if one corrects for the attenuated policy rate response at the lower frequency. Hence, both the shapes and the relative magnitudes of the responses across different frequencies are consistent.

Estimating impulse responses based on annual data yields accurate results in our setting. In principle, the approach of aggregating shocks to a lower frequency could also be applied to other environments in which detailed micro data is available only at an annual frequency. However, time aggregation can introduce bias and its effects should therefore always be carefully investigated. In Appendix A.10, we provide further intuition on the type of impulse responses that are more likely to result in similar responses across various frequencies. Apart from the aggregate responses, our interest lies in differences in the reactions to monetary policy in the cross section of households when using the micro data. Potential attenuation arising from time aggregation of the shocks therefore plays a smaller role for that part of the analysis.

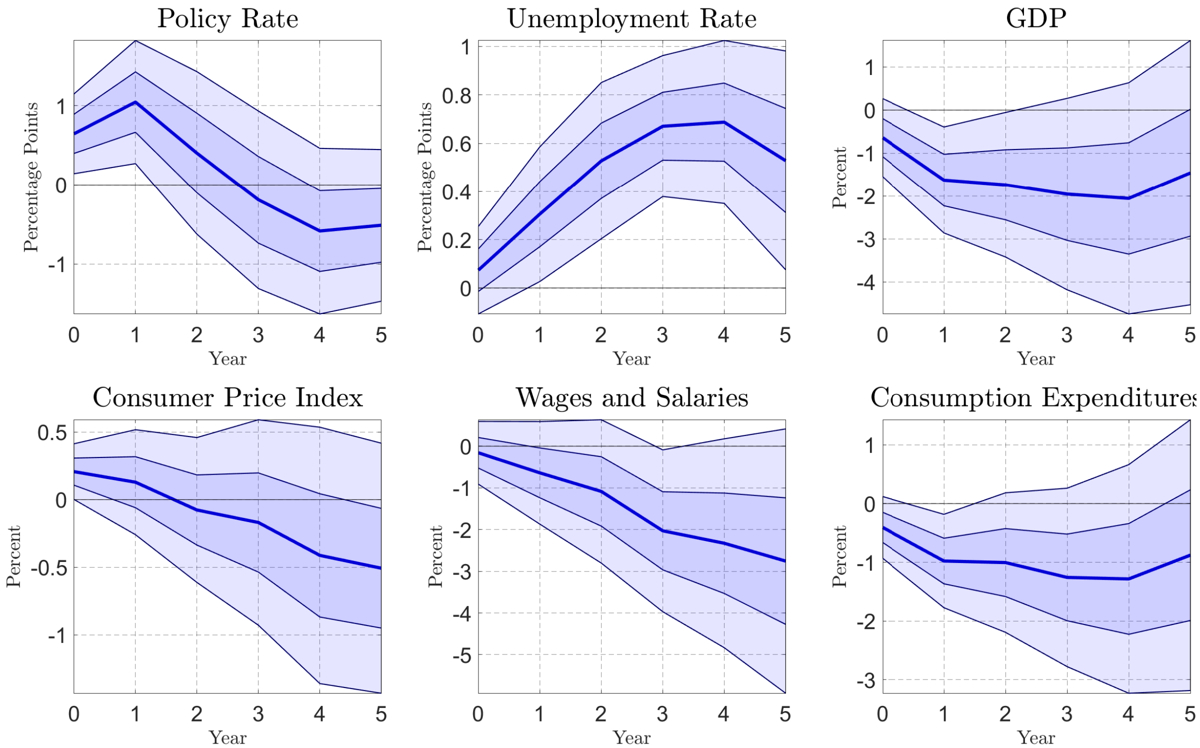

Figure 3: **Impulse Responses at an Annual Frequency.**
Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at an annual frequency, based on the local projection approach in (2). 95 and 68 percent confidence bands shown, using Newey and West (1987) standard errors. Additional impulse responses at an annual frequency are shown in Figures A.10 and A.11 in Appendix A.7.

3 Administrative Data

At the micro level, we base our study on Norwegian administrative data. Norway levies both income and wealth taxes on its inhabitants. The tax authority therefore collects information on all sources of income and balance sheet components down to detailed asset categories. The data are largely third-party reported, meaning that employers and banks report income and balance sheet information directly to the tax authorities. Below, we describe the data in detail, including the different sources, the procedure we follow to impute consumption, and the sample restrictions we impose, before turning to descriptive statistics.

3.1 Data Sources

We combine a number of Norwegian administrative registries maintained by Statistics Norway. All registries contain unique identifiers at the individual and household level, allowing us to link information from multiple sources. Our unit of observation is the household since saving and consumption decisions are made at the household level, and because wealth is taxed at the household level. We combine a rich longitudinal database covering every resident (containing socioeconomic information including sex, age, marital status, family links, educational attainment, and geographical identifiers), the individual tax registry, the Norwegian shareholder registry on listed and unlisted stock holdings, balance sheet data for listed and unlisted companies, and registries of housing transactions and ownership. All income flows are annual and assets are valued at the end of the year.

For our study, the Norwegian data have several advantages. First, the balance sheets and income statements of households are observed across multiple time periods. We are therefore able to construct a panel and follow the responses of households to monetary policy shocks across multiple years. Second, our data cover the universe of Norwegian households, allowing us to investigate the responses to monetary policy across many dimensions without running into issues of small samples. Third, the data are not top-coded and the only sources of attrition are death and migration. Fourth, the data are third-party reported with the exception of a few smaller items in the tax return, limiting the scope for measurement error.[13]

3.2 Imputed Consumption Expenditures

We compute a measure of consumption expenditures for each household using the budget constraint:[14]

$$c_{i,t} = inc_{i,t} - s_{i,t} \quad , \tag{3}$$

where $c_{i,t}$ is consumption for household i in year t, $inc_{i,t}$ is disposable income, and $s_{i,t}$ is saving, defined as the change in net wealth excluding capital gains. Disposable income, $inc_{i,t}$, is observed in the data as the sum of labor income (salary and business income), capital income (dividends and interest income net of interest expenses), transfers (pensions, social security, and unemployment insurance), and retained earnings in private businesses, net of taxes. Net wealth is the sum of all assets (stocks, bonds, stock funds, private business, deposits, housing, vehicles, and outstanding claims) minus liabilities (mortgage, consumer, and student debt).

The main challenge of consumption imputation is to compute the relevant measure of saving excluding capital gains.[15] A measure of saving *including* capital gains is directly observed because information on net wealth is available for the beginning and end of each year. To arrive at a measure of saving *excluding* capital gains, estimates of (realized and unrealized) capital gains are needed. Four types of assets accrue capital gains in our data (housing, stocks, stock funds, and private business) and different methods for calculating capital gains are applied for each of them. For housing, we observe transactions and compute capital gains as the change in housing wealth that is not due to housing transactions.[16] For stocks, we use the stockholder registry and capital

[13] The most important (partially) self-reported positions are outstanding claims, inheritance, and foreign assets.

[14] Consumption imputation has by now been widely applied in the literature (see, e.g. Leth-Petersen 2010; Fagereng, Holm and Natvik 2021; Eika, Mogstad and Vestad 2020). Fagereng and Halvorsen (2017) provide details based on Norwegian data, although our method differs from theirs because we also utilize detailed transaction data on stocks and housing.

[15] Baker, Kueng, Meyer and Pagel (2018) compare imputed consumption with transaction-level data and show that systematic measurement errors could be linked to households with large holdings of assets that experience capital gains. We utilize detailed transaction data from stocks and housing to limit the extent of these measurement errors and test the robustness of our main results to the consumption imputation in Section 4.2.4 and Appendix C.3.

[16] Our measure of housing wealth is from Fagereng, Holm and Torstensen (2020). They combine detailed transaction

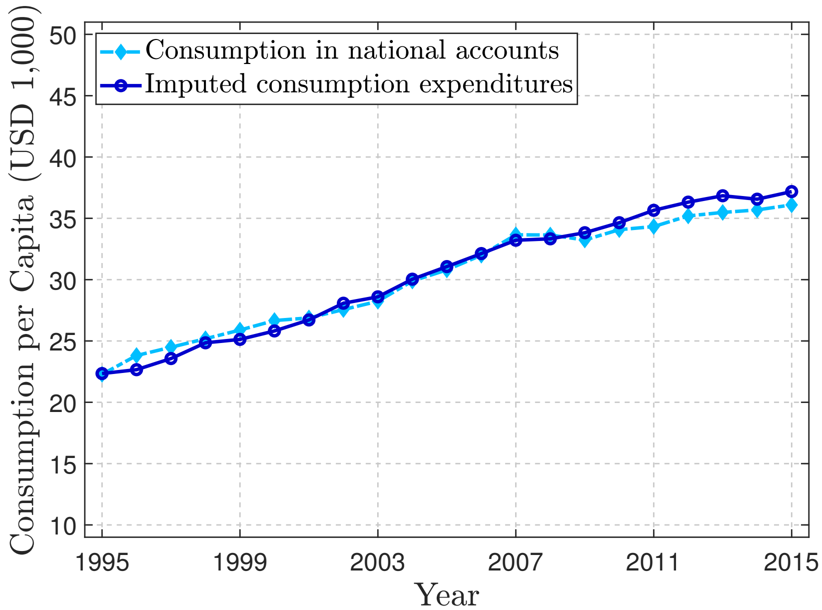

Figure 4: **Imputed Consumption Expenditures and Consumption in the National Accounts.**
Notes: Shown are real imputed consumption expenditures based on equation (3) and real consumption in national accounts. To make the series comparable, we exclude imputed rents for owner-occupied housing from the national accounts. All values are in real U.S. dollars.

gains on individual stocks after 2005 and average capital gains for stocks traded on the Oslo Stock Exchange prior to 2005. For stock funds, we use the measure of capital gains on stock funds from the national accounts. And for private businesses, we assume that capital gains are zero unless we observe that the company holds listed stocks on its balance sheet. If a firm holds stocks, we attribute its share of capital gains on the stocks to the owner of the private business according to the individual's ownership share in the firm. Appendix B.1 presents more details on how we compute capital gains. Figure 4 compares the imputed consumption expenditures with consumption in the national accounts. The two consumption series follow each other closely.

3.3 Sample Restrictions

We impose some minor sample restrictions. First, we focus on the adult population older than 20 years. Second, since our measure of consumption applies at the household level, we drop household-year observations in which individuals change marital status between single and couple. Third, we only include households that reside in Norway in at least two consecutive years. Fourth, we require individuals to have income and consumption above the minimum level in the Norwegian social security scheme.[17] And finally, since there may be assets that experience sharp revaluations or assets that do not appear on the balance sheet in some years, we require the growth rate of imputed consumption expenditures to be less than 50 percent in absolute value.[18] Table 2

data and information on housing unit characteristics to estimate housing wealth at the household-level using machine learning methods.

[17]The minimum level in the Norwegian social security was approximately USD 11,000 in 2015.

[18]After we impose the bound on age, the sample restrictions eliminate the following fraction of observations: change in marital status (4 percent), two consecutive tax filings in Norway (5 percent), disposable income and consumption

Table 2: **Summary Statistics.**

	Mean	SD	P10	Median	P90
Age	51.63	17.85	28.00	50.00	77.00
Consumption	43,091	159,368	22,099	37,714	65,424
Disposable income	43,437	81,284	23,616	39,833	63,817
Income before tax	58,827	89,245	26,940	52,875	93,096
Labor income	44,210	42,362	0	43,977	92,636
Net capital income	-1,692	21,031	-8,263	-892	2,355
Dividend income	429	19,841	0	0	15
Interest income	873	3,150	5	198	2,207
Interest expenses	3,316	5,072	0	1,631	8,970
Total assets	371,601	1,292,982	5,588	281,798	782,215
Liquid assets	31,337	75,379	565	11,262	78,912
Deposits	26,569	59,632	465	9,065	67,554
Bonds	1,015	13,660	0	0	0
Risky financial assets	4,261	293,320	0	0	8,038
Stocks	1,945	292,750	0	0	660
Stock funds	2,316	12,507	0	0	5,339
Housing	321,580	371,837	0	248,128	703,170
Total debt	73,658	885,968	0	33,954	186,687
Observations per year	1,909,603	83,648	1,821,377	1,864,722	2,032,543

Notes: The table shows summary statistics for the estimation sample. Disposable income is the sum of labor income, capital income, transfers, and retained earnings in private businesses, net of taxes. Liquid assets is the sum of deposits, bonds, stocks held directly, and stock funds. Risky financial assets consist of stocks and stock funds. Stocks also includes stocks held indirectly through holding companies. Total debt includes mortgages, consumer debt, and student debt. All values except age are in U.S. dollars, 2011 prices.

presents summary statistics for the estimation sample.

3.4 Institutional Setting

In Norway, mortgage and deposit rates of existing contracts are sensitive to policy rate changes. The standard mortgage contract features an adjustable rate, accounting for more than 90 percent of all outstanding mortgage debt, in contrast to the United States where such contracts are typically issued with a fixed rate.[19] In the deposit market, demand deposits account for more than 85 percent of all deposit contracts and deposits are the predominant form of liquid asset holdings, as can be seen in Table 2.[20] Figure B.1 in Appendix B.2 demonstrates that deposit rates closely track the policy rate. In the United States, the rates paid on demand deposits are generally more sticky, resulting in substantial changes in the deposit spread over the monetary policy cycle (see, e.g., Drechsler, Savov and Schnabl, 2017).

above the social security minimum (6 percent), and bounds on consumption growth (19 percent).

[19] In the years, in which data on aggregate loans to households are available (2013 - 2018), between 90 and 93 percent of all outstanding debt issued to households took the form of adjustable mortgage contracts.

[20] Between 1996 and 2015, the share of deposits with adjustable rates was between 85 and 92 percent.

A potential concern is that the data do not cover information on pension wealth. However, the absence of this information is unlikely to affect our results for the following reason. The large majority of pension wealth held by Norwegian households is made up of defined benefits associated with the public pension system. As households work, they accumulate pension points that are translated into benefit payments when they retire. The pension income they receive, which we record as transfers, is independent of monetary policy and the performance of financial markets more broadly. All risk is instead born by the Norwegian government, which holds substantial wealth in the Norwegian Government Pension Fund Global ("the oil fund").

3.5 Income, Wealth, and Liquid Assets

Following the recent literature on HANK models, we initially focus on the effects of monetary policy on households located at different segments of the liquid asset distribution. Below, a household's liquid assets are given by the sum of its portfolio positions in bank deposits, government debt, corporate bonds, publicly traded stocks, and stock fund shares.[21] The remainder of the section illustrates characteristics of the wealth and income composition of households along the liquid asset distribution.

Asset and debt holdings are concentrated at opposite ends of the distribution. Figure 5 shows the cumulative shares of the most significant asset and liability classes. While assets, particularly financial assets, are concentrated among households with larger liquid asset holdings, those with smaller liquid asset positions hold a disproportionately large share of debt. The value fractions of deposits, stocks including mutual funds, and bonds that lie in the hands of the bottom half of the distribution are 6.7, 4.2, and 2.6 percent, respectively. Conversely, the top 10 percent of the distribution hold 50.9 percent of all deposits, 59.9 percent of all stocks, and 61.1 percent of all bonds contained in household portfolios. Illiquid assets in the form of housing are less concentrated along the liquid asset distribution, although housing wealth is equally more prevalent among households with high liquid assets. The opposite is true for debt, with consumer debt being more unequally distributed than mortgages.

The concentration of assets and debt is reflected in the composition of household portfolios and income. The left panel of Figure 6 illustrates the average portfolio shares of financial assets and housing as well as the average ratio of total debt to total assets in groups of households ordered by liquid asset holdings. The right panel shows financial and nonfinancial income as fractions of disposable income.[22] Each group contains 5 percent of all observations, except at the upper end of the distribution, where we separate out observations located between the 95th and 99th, between the 99th and 99.9th, and above the 99.9th percentile of the liquid asset distribution.

[21] This definition follows Kaplan, Violante and Weidner (2014). The results discussed in Section 4 are robust to excluding equity and mutual fund shares from liquid assets.

[22] The shares of further asset classes like physical nonhousing wealth are omitted here for clarity. Net financial income is given by capital revenues net of capital expenditures. Nonfinancial income is the sum of labor income and transfers net of taxes.

At the bottom and in the middle of the distribution, most wealth is held in the form of housing. Debt exceeds total assets below the 35th percentile and financial assets below the 70th percentile. This implies that, on average, net financial income is negative and hence that disposable income is smaller than nonfinancial income in the majority of groups. At the top of the distribution, financial asset holdings gain while housing and debt decline in importance. For households in the 99-99.9 percent group and in the top 0.1 percent group, the value of the financial portfolio exceeds that of real estate on average. Correspondingly, the ratio of net financial income to disposable income increases from -6.9 percent in the group containing the median household, which is negative as households are net debtors, to more than 55 percent in the top 0.1 percent group.

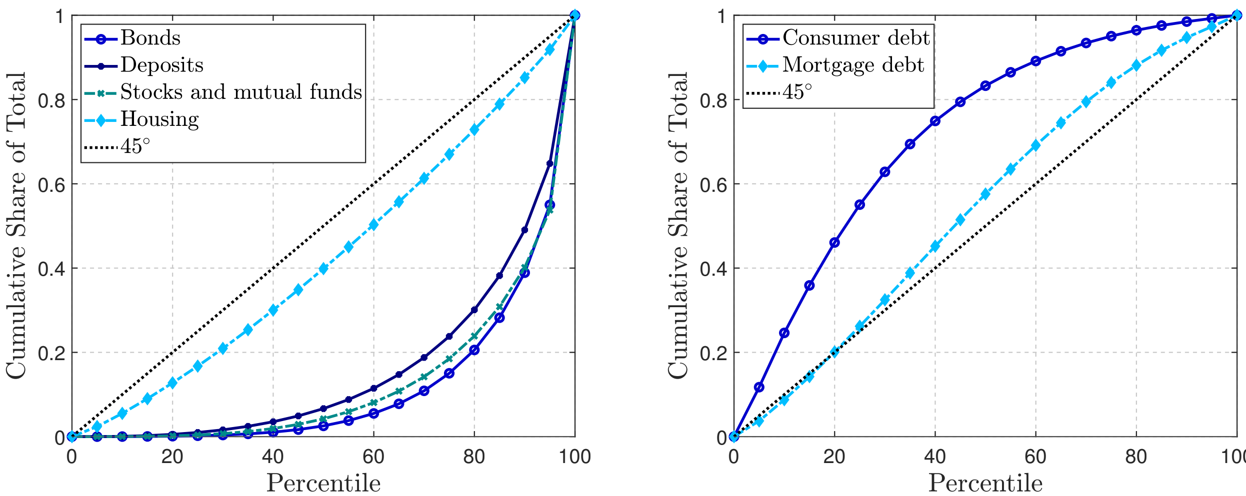

Figure 5: **Cumulative Shares of Assets and Liabilities.**
Notes: Cumulative shares by asset and debt type against percentiles of liquid asset distribution.

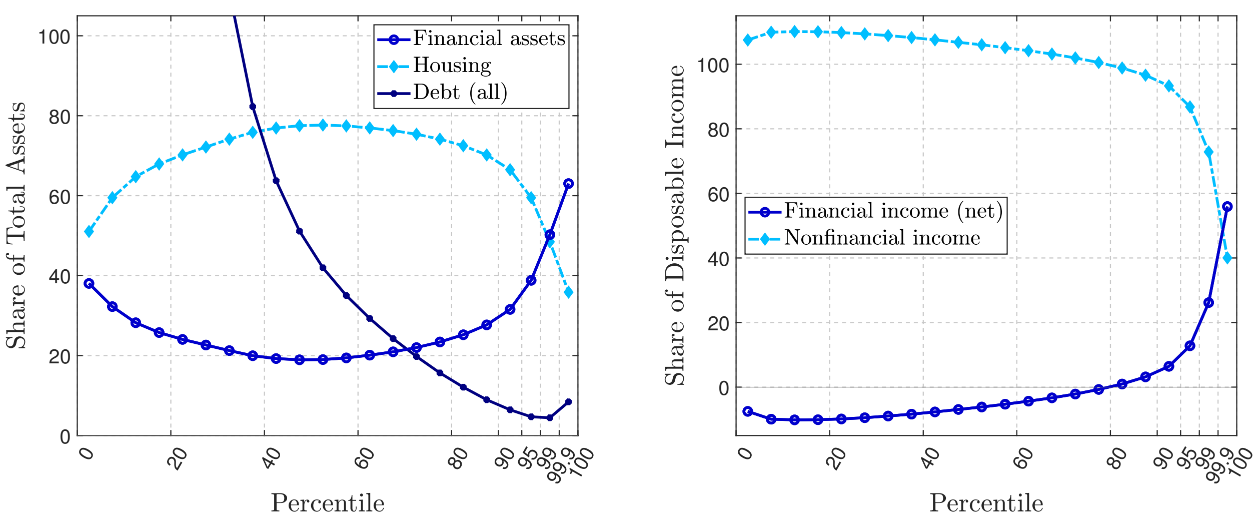

Figure 6: **Portfolio and Cash-Flow Shares.**
Notes: Shares of total assets (left panel) and disposable income (right panel) against percentiles of liquid asset distribution. Group averages shown.

4 Monetary Policy Transmission at the Household Level

In this section, we use the shocks derived in Section 2 together with the data set described in the previous section to investigate the transmission of monetary policy at the household level. The section first shows estimates of the aggregate effects of monetary policy implied by the household-level data. It then disaggregates the results exploring the mechanisms that underlie them. We focus on the role played by the liquid asset positions of households to draw conclusions about the channels emphasized by the HANK literature and on interest rate exposure to assess the significance of cash-flow effects.

4.1 Micro-Macro Responses

As a first pass, we use the whole sample and obtain "macro impulse responses based on micro data." To this end, we estimate local projections of the form

$$\frac{y_{i,t+h} - y_{i,t-1}}{\overline{y}_{t-1}} = \delta_i^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{k=1}^{K} \gamma_k^h X_{i,t-k} + u_{i,t}^h \quad , \tag{4}$$

where $y_{i,t}$ is some outcome variable of interest specific to household i at time t, now denoted in (real) levels (e.g., household disposable income at constant prices). Further, δ_i^h denotes a household-specific constant for horizon h, and $X_{i,t}$ is a vector of controls. Following the specification in (2), we include three years of lagged values of ϵ_t^{MP} and add two years of lagged one-year growth rates of the dependent variable.[23] Notably, the dependent variable is given by the change in y_i from $t-1$ to $t+h$, normalized by the average value $\overline{y}_{t-1}$ across all households at time $t-1$. This normalization makes β^h comparable to the corresponding coefficient estimated on aggregate data in Section 2.2.[24]

The estimated impulse responses are shown in Figure 7 for a selection of household variables. After a monetary tightening, consumption expenditures and disposable income fall. The decline of disposable income can be separated into the reactions of nonfinancial income (earned income and net transfers) and financial income (capital revenues minus capital expenditures). The former mimics the disposable income response and the latter is separated into its two subcomponents. Both capital revenues and expenditures follow the response of the policy rate in Figure 3. Initially they increase and subsequently fall. Their response suggests that they are driven by movements in interest income and expenditures, channels that we investigate in more detail below. Further, household debt, wealth, and risky financial assets fall, whereas the response of safe assets

[23] Specifically, we include $\left(y_{i,t-1} - y_{i,t-2}\right)/\overline{y}_{t-2}$ and $\left(y_{i,t-2} - y_{i,t-3}\right)/\overline{y}_{t-3}$ as controls. To choose the lag length for the dependent variable, we again consult the Akaike and Bayesian information criteria. Across various outcome variables and impulse response horizons, the chosen lag length of two years is again a reasonable compromise. We tested (and confirmed) the robustness of our main results to the choice of the lag length for the dependent variable and the shock.

[24] In contrast, when normalizing by $y_{i,t-1}$ instead, β^h would give the average of the household-specific percentage changes to a monetary policy shock. While an interesting estimate in its own right, the interpretation of β^h would differ from the one based on macro aggregate data and may be driven by extreme observations of some households.

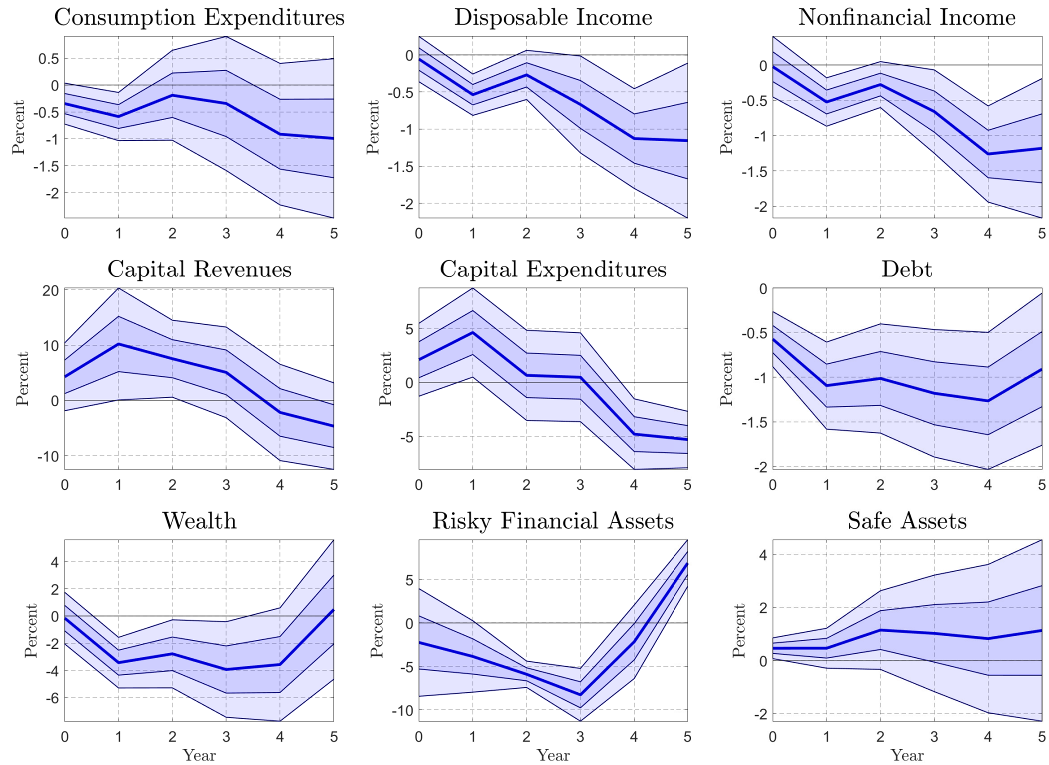

Figure 7: **Micro-Macro Impulse Responses.**
Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at an annual frequency, based on the local projection approach in (4). 95 and 68 percent confidence bands shown, using Driscoll and Kraay (1998) standard errors.

is slightly positive.[25] Overall, the responses go in the expected directions and they are statistically different from zero at the 95 percent or at least at the 68 percent confidence interval.[26]

4.2 Impulse Responses along the Liquid Asset Distribution

To investigate the role of liquid asset holdings in monetary policy transmission, we divide households into groups of equal size indexed by $g = 1, 2, ..., 10$ and estimate separate impulse responses for each group. A household i is allocated to group g in period t if its liquid asset holdings in $t-1$ fell between the $(g-1)$-th and g-th decile of the distribution. Ordering households according to lagged liquid asset holdings guarantees that the group allocation is not influenced by the shock

[25]Household debt consists of mortgage, consumer, and student debt. Household wealth is given by housing wealth, financial net worth, the value of vehicles (boats and cars), and outstanding claims (private loans and receivables). We find that the response of total household wealth is almost entirely driven by housing wealth, which is by far the largest component (see also Figure 4 in Fagereng, Holm, Moll and Natvik, 2019). Risky financial assets consist of stock holdings and mutual stock fund holdings. The inclusion of the latter is due to the fact that mutual stock funds typically invest in risky assets in Norway. Safe assets are given by deposits, bond holdings, and outstanding claims.

[26]The micro data also allow us to compute measures of inequality and their responses to monetary policy. We focus on disposable income, consumption expenditures, and wealth. For these three variables, we calculate the difference between the 90th and the 10th percentile of log levels of each distribution (p90-p10). Figure C.1 in Appendix C.1 shows the results based on the local projection approach in (2). Income and consumption inequality increase substantially after four to five years, confirming the findings by Coibion et al. (2017) based on U.S. data. In addition, we find that wealth inequality decreases in response to a monetary policy tightening, which is linked to the decline in asset prices associated with a monetary policy tightening (see Figure 14).

occurring in period t. For each group g, the setup of the local projections is

$$\frac{y_{i,t+h} - y_{i,t-1}}{inc_{i,t-1}} = \delta_i^h + \beta_g^h \cdot \epsilon_t^{MP} + \sum_{k=1}^{K} \gamma_{g,k}^h X_{i,t-k} + u_{i,t}^h \qquad \forall i \in g \ , \tag{5}$$

where the notation is as in Section 4.1 and $h = 0, 1, ..., 5$. As in (4), $X_{i,t}$ includes three years of lagged values of ϵ_t^{MP} and two years of lagged one-year growth rates of the dependent variable. The main difference to (4), apart from the grouping of households, lies in the dependent variable. Changes of y_i from $t-1$ to $t+h$ are expressed in $t-1$ units of disposable income $inc_{i,t-1}$.[27] This normalization ensures that the estimated coefficients β_g^h are comparable across different variables and that they give an indication of the economic importance of the effects for the households.

Before turning to the results, we outline the effects of monetary policy on household consumption predicted by theory to fix ideas and introduce the terminology used below. First, a change in the interest rate affects the price of consumption tomorrow relative to consumption today resulting in a substitution effect. Second, an interest rate adjustment gives rise to a standard income effect, which provides an incentive for savers to increase and for borrowers to decrease their consumption in response to a rise in the interest rate. Third, if households are particularly sensitive to transitory income changes (e.g., because they are close to a borrowing constraint or subject to information frictions), their consumption may exhibit an elevated positive response to their net interest income. We refer to the effect of changes in interest income or expenses on consumption in excess of that implied by the standard income and substitution effect as a "cash-flow effect." Fourth, monetary policy may trigger general equilibrium price responses, which themselves can affect consumption through income, substitution, and cash-flow effects, analogous to those above. For example, a consumption decline brought about by a rise in the interest rate that lowers labor earnings through the wage rate constitutes an indirect income effect of this type. Finally, interest rate changes may lead to a revaluation of financial and nonfinancial assets contained in household portfolios, which can be transmitted to consumption through their effect on household wealth or collateral, for example. We categorize these effects as consumption due to capital gains or losses.[28]

4.2.1 Consumption, Income, and Saving

Figure 8 shows the impulse responses of consumption and the ratio of consumption to disposable income to a positive 1 percentage point shock. The corresponding responses of disposable income and saving are shown in Figure 9. As before, saving is an "active" flow, that is, it is the part of the change in a household's assets which is not due to capital gains or losses.[29] For each variable, the figures portray the complete impulse responses of all ten groups in the top row and a time-averaged version which facilitates a two-dimensional representation in the bottom row. Since the figures display the impulse responses across the entire distribution, they only show the

[27] Changes in the consumption-to-income ratio are not normalized by lagged income.

[28] An example is a consumption adjustment in response to stock price movements that are not merely a compensation for changes in dividends.

[29] Throughout, we distinguish between *saving* (flow) and *savings* (stock).

point estimates, not the corresponding confidence intervals, to allow for a clear presentation. For each figure contained in this section, we include two sets of results in Appendix C.2. Figures C.2-C.3 show the full impulse responses and the associated confidence intervals for selected groups ($g = 1, 5, 10$) and Figures C.4-C.5 the relative responses across groups, which allow to evaluate the significance of group differences.

The responses to a monetary policy shock differ systematically according to the level of liquid asset holdings. For the bottom 10 percent of households, disposable income stays nearly unchanged initially and then gradually falls. The decline is significant at the 5 percent level from year 3 onward and, on average, amounts to about 1.4 percent of disposable income in years 4 to 5 after the shock. The responses of saving and the consumption-to-income ratio are small in comparison, so that consumption closely follows income and also starts to fall statistically significantly after three years. Consumption of households located close to the median of the distribution is less responsive. In conjunction with their income decline in years 2 and 3, they allow their consumption-to-income ratio to rise and hence saving to fall. Both responses become statistically different from zero after the second year. In years 4 and 5, consumption is lowered but, relative to initial household income, the size of the decline is smaller than for the bottom 10 percent of households. The differences in the consumption response between these groups are equally significant at the 5 percent level.

At the top of the distribution, the responses are markedly different from those at the median and the bottom. Disposable income in the top group increases by about 1.4 percent on average in years 0 and 1. Saving rises by around 0.7 percent of household income, implying that consumption increases by roughly the same amount. The positive consumption response is significant at the 5 percent level in year 1. The ensuing sharp fall in income is associated with a statistically significant consumption adjustment of about the same size four to five years after the shock has occurred. At this impulse response horizon, the consumption response is of similar magnitude as that at the bottom of the distribution. As a result, the peak consumption response is inverse U-shaped along the liquid asset distribution.

Several of the results are worth stressing. At the bottom, the peak consumption response is significantly larger than around the median and the response of the consumption-to-income ratio is comparably small when income begins to fall. Both of these observations are consistent with the HANK literature, although the mechanisms at work may differ. The two-asset model by Kaplan, Moll and Violante (2018) predicts that households with small liquid asset holdings are close to a kink in their budget constraint and therefore have large MPCs, implying that monetary policy affects their consumption particularly strongly and the consumption-to-income ratio only weakly. Our estimates show that the disposable income of households at the bottom begins to fall in years 2 to 3 and further declines in years 4 to 5, while saving remains relatively stable. The lack of a strong saving response is consistent with hand-to-mouth behavior at these horizons. Since saving increases initially and consumption falls somewhat when the policy rate rises, the estimates do not

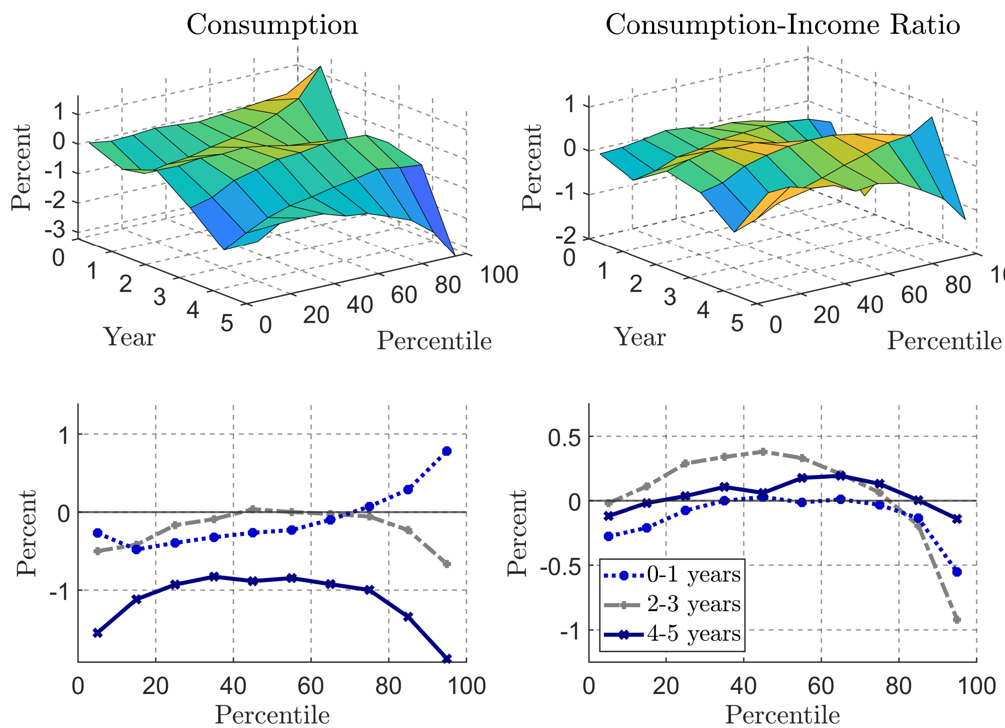

Figure 8: **Impulse Responses of Consumption and Consumption-Income Ratio.**
Notes: Changes relative to lagged net income (left) and changes of ratio (right) at segments of the liquid asset distribution. See Figures C.2-C.5 in Appendix C.2 for group responses and group comparisons with confidence bands.

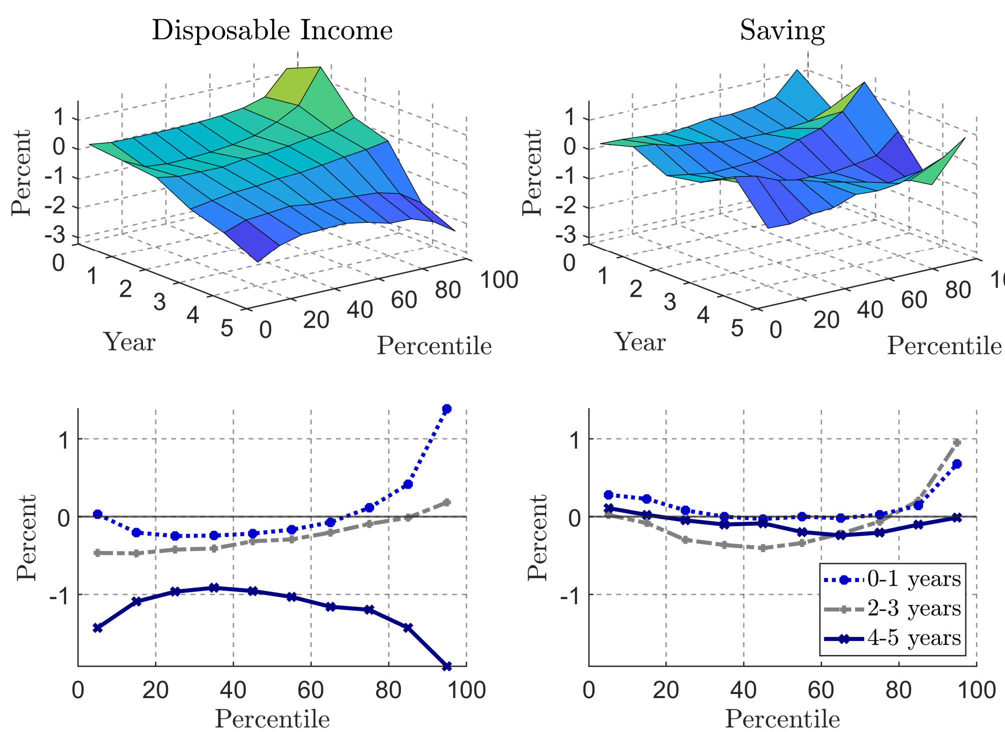

Figure 9: **Impulse Responses of Net Income and Saving.**
Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures C.2-C.5 in Appendix C.2 for group responses and group comparisons with confidence bands.

allow us to reject that intertemporal substitution may play a role on impact. Around the median of the distribution, the decline in income after two to three years is matched by dis-saving. Hence, households in the middle of the distribution show signs of consumption-smoothing behavior, although they equally let income changes feed into consumption at the 4 to 5 year horizon.[30]

Our findings for the top of the distribution are inconsistent with the predictions of canonical HANK models. In these models, it is possible that households with high amounts of liquid asset holdings show a large and positive consumption response to a rise in the policy rate if the standard income effect dominates the substitution effect of the interest rate change. However, the magnitude of the consumption response relative to the income response that we observe is a lot larger than what can be generated by current HANK models. Put differently, for an income response of realistic size, structural models are unable to give rise to a consumption response that matches our estimates. The reason is that high-liquidity households in HANK models smooth the effect of temporary income changes on consumption closely in line with the permanent income hypothesis, implying that they have very low MPCs. We illustrate this point using a structural model in Appendix D. In the model, the response of consumption is too small relative to the response of disposable income even when we calibrate it such that high-liquidity households hold significantly more liquid assets than in the data and the income effect is permitted to dominate the substitution effect, as shown in Figure D.1.

Instead, our results are consistent with high MPCs even for high-liquidity households and hence with cash-flow effects playing an important role. In line with this view, Fagereng, Holm and Natvik (2021) find MPCs out of lottery prizes that are substantially larger than zero for households with substantial levels of liquidity. Based on an extended sample that matches our liquidity groups, we revisit their estimates and plot annual MPCs for the contemporaneous year in Figure 10. While the estimated MPC is declining with liquid asset holdings, it is still sizable for the top ten percent of the distribution (around 0.42), in accordance with our findings.[31] In the following sections, we investigate the effect of monetary policy on household income and consumption in more detail.

[30] The seemingly high MPC on impact may be a result of two motives present at the same time. Households wish to increase saving and reduce consumption in response to the interest rate hike and to lower saving due to the decline in disposable income.

[31] In Figure E.3 in Appendix E.5, we apply these estimates, together with the ones for year 1, to the estimated disposable income response in Figure 9 and compare the implied consumption response to the estimated impact consumption response shown in Figure 8. Especially for the top of the distribution, the two estimates line up closely.

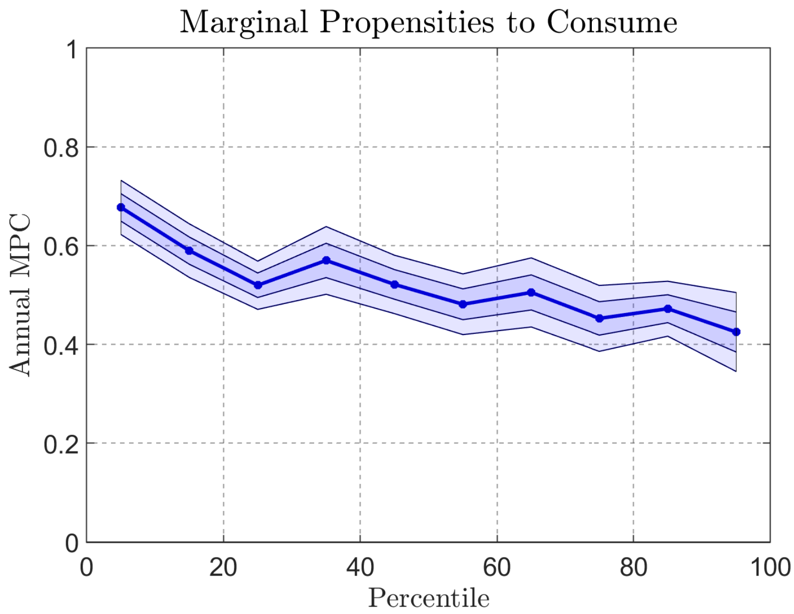

Figure 10: **Estimated Annual MPCs across the Liquid Asset Distribution.**
Notes: The figure shows annual MPCs for the contemporaneous year across the liquid asset distribution, with 68 and 95 percent confidence bands. The estimates are obtained using lottery winnings following Fagereng, Holm and Natvik (2021) based on an extended sample that matches the different liquidity groups.

4.2.2 Decomposing the Income Response

Below, we decompose the income responses portrayed in Figure 9 into various subcomponents. A first set of results is contained in Figure 11. It shows the impulse responses of the two most important income components, net financial income, defined as the difference between capital revenues and capital expenditures, and nonfinancial income, given by the sum of labor earnings, pensions, and other transfers net of taxes.

While nonfinancial income shifts household income more than financial income at the bottom of the liquid asset distribution, the opposite is true for households at the top. In the bottom group, nonfinancial income begins to gradually fall in the 2 to 3 year interval until an average decline of about 2.1 percent of lagged disposable income is reached in years 4 and 5. Net financial income is lowered by 0.3 percent in the first two years but elevated by 0.6 percent in the final two years. For the top group, the eventual decline in nonfinancial income amounts to about 1.3 percent. Importantly, net financial income initially increases by 2.3 percent of disposable income before a decline of 1.6 percent is reached. All of these estimates are significant at the 5 percent level.

The differences in the responses along the liquid asset distribution are related to the patterns observed in the composition of income illustrated in Section 3.5. Figure 6 shows that the average labor income share is significantly larger at the bottom of the distribution than at the top. In contrast, the average net financial income share (in absolute value) is larger at the top than at the

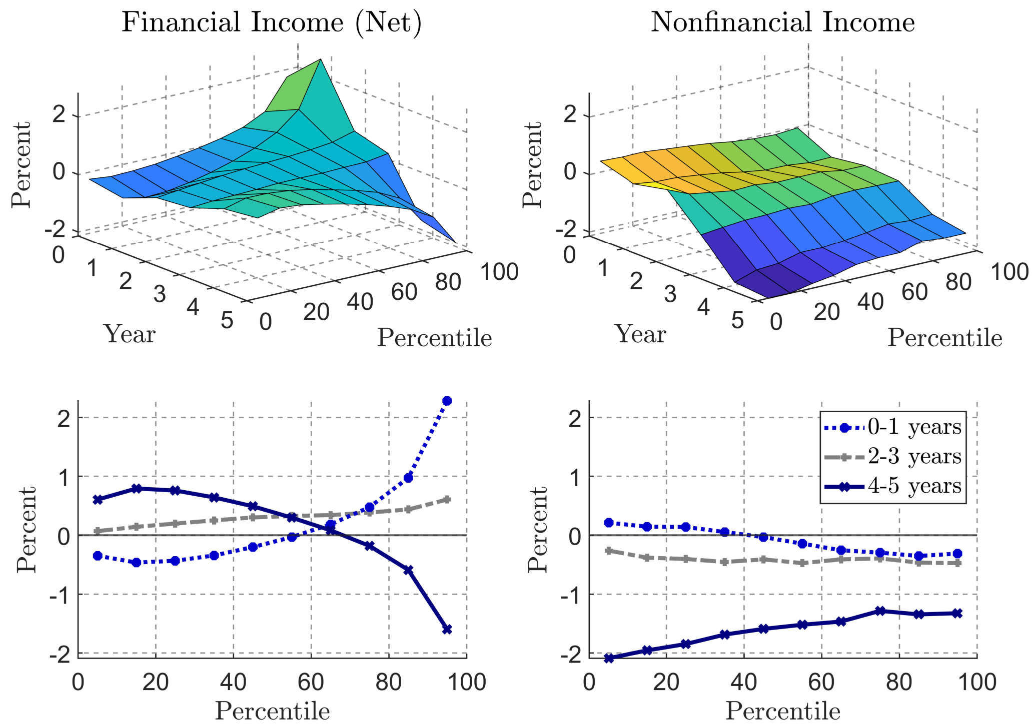

Figure 11: **Impulse Responses of Net Financial Income and Nonfinancial Income.**
Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures C.2-C.5 in Appendix C.2 for group responses and group comparisons with confidence bands.

bottom. Correspondingly, the nonfinancial income response is particularly strong for households with small liquid asset holdings, while the net financial income response is more pronounced for households with large liquid asset holdings. The initial increase in the interest rate following the shock lets financial income rise for households at the top and decline for households at the bottom, in line with average net financial income being of opposite sign. At both ends, the response in the last two years nearly mirrors that in the first two years as the policy rate drops below its initial level in year 3 after the shock.

The response of net financial income is decomposed further into the responses of capital revenues and capital expenditures in Figure 12. The estimates reflect the marked concentration of creditors at the upper end and debtors at the lower end of the liquid asset distribution. Both capital revenues and capital expenditures closely follow the policy rate. The latter is explained by the high prevalence of adjustable-rate mortgages in Norway. To investigate the former, we break the capital revenue response down further. Figure 13 shows the impulse responses of interest income and dividend income. Note that the dividend income responses are portrayed on a magnified scale. Nearly the entire capital income response is due to changes in interest income. Dividend income changes do not affect household income in a quantitatively meaningful way following monetary policy shocks. A likely explanation is corporate dividend smoothing, a well-documented empirical finding that goes back at least to Lintner (1956).

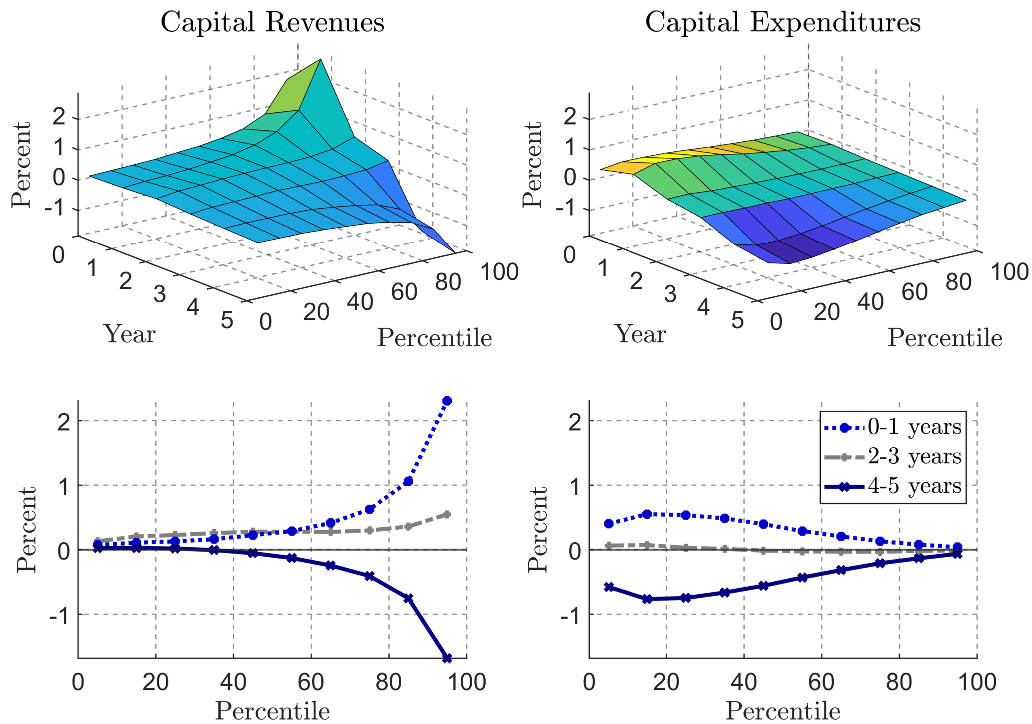

Figure 12: **Impulse Responses of Capital Revenues and Expenditures.**
Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures C.2-C.5 in Appendix C.2 for group responses and group comparisons with confidence bands.

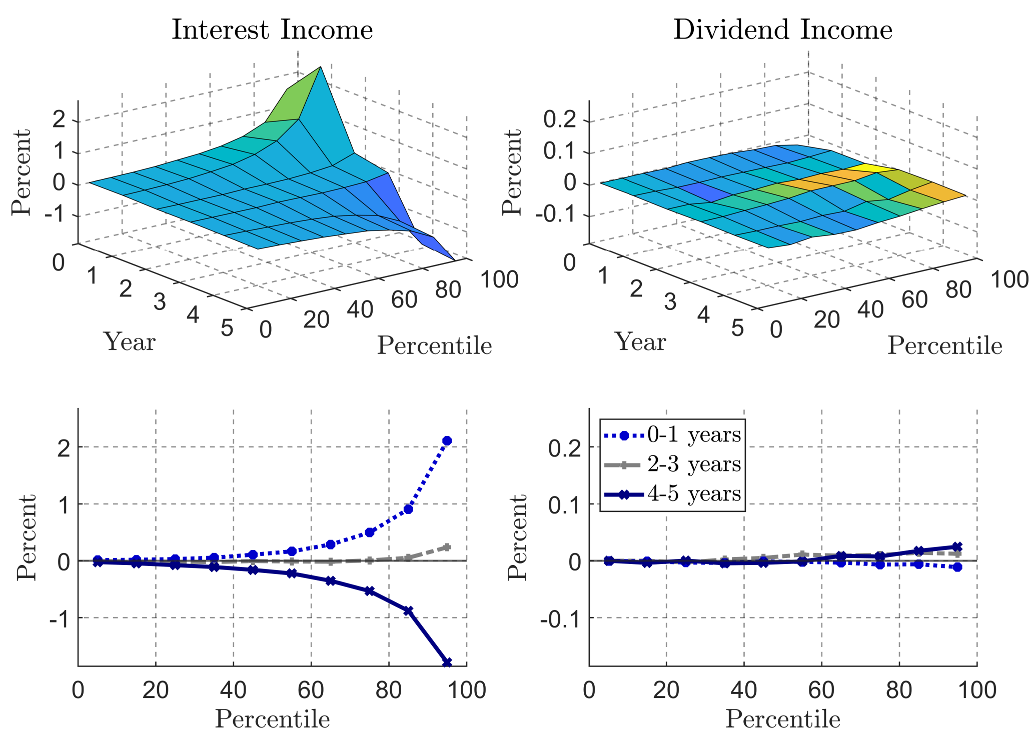

Figure 13: **Impulse Responses of Interest and Dividend Income.**
Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures C.2-C.5 in Appendix C.2 for group responses and group comparisons with confidence bands.

Figures 11-13 show that there are slow-responding and fast-responding income components. Together with the estimates in the previous section, which suggest that household consumption is responsive to the income changes induced by monetary policy, this implies that the consumption response to a policy rate adjustment can be divided into different stages. Financial income, driven by interest income and expenses, responds on impact. Consequently, cash-flow effects play an important role in the early stages of the response. They are of opposite sign at both ends of the distribution, explaining the positive consumption response at the top and the negative response at the bottom. The nonfinancial income response takes time to build up, but is ultimately sizable and negative across the entire distribution, which suggests that indirect effects dominate the later stages of the response. A possible way to rationalize a delayed nonfinancial income response accompanied by a concurrent consumption contraction is a Keynesian consumption-investment feedback as described in Auclert, Rognlie and Straub (2020) and Bilbiie, Känzig and Surico (2020). In the HANK model by Auclert, Rognlie and Straub, a rise in the interest rate lowers investment leading to a decline in aggregate demand and wages. The fall in labor income in turn reduces consumption and thus output demand further setting off a multiplier effect. Adjustment costs imply that investment reacts sluggishly and stickiness in household expectations implies that consumption responds at the same time as labor income.[32] Bilbiie, Känzig and Surico (2020) show that a multiplier of this type is amplified if the income of households with particularly high MPCs responds disproportionately strongly.

4.2.3 Capital Gains

To ascertain that households do not merely respond to changes in the valuation of their assets, we plot the impulse responses of capital gains in housing and risky financial assets in Figure 14. Since capital gains fall substantially in years 0 to 1, they cannot explain the large initial consumption expansion at the top of the liquid asset distribution. This finding is consistent with the low MPCs typically estimated out of capital gains (Chodorow-Reich, Nenov and Simsek, 2021; Fagereng et al., 2019; Di Maggio, Kermani and Majlesi, 2020). The estimates reflect that housing is more equally distributed than risky financial asset holdings and that housing units on average represent a significantly larger share of household portfolios in all groups considered. The impulse responses of capital gains in housing are therefore relatively uniform across the distribution and larger in absolute value despite stock prices being more sensitive to monetary policy than house prices.

4.2.4 Robustness

The results are robust across a range of different specifications. The corresponding estimates are relegated to Appendix C.3. Figure C.6 demonstrates that the patterns in the consumption responses are equally pronounced if the consumption growth rate is considered. The same is true if we include a measure of housing service flows from owner-occupied housing following Fagereng,

[32] Their formulation of sticky information in the spirit of Gabaix and Laibson (2001) and Mankiw and Reis (2002) is silent about the sources of the information friction. However, D'Acunto et al. (2019) estimate that high-IQ men in Finland are more likely to update their information about prevailing interest rates than low-IQ men.

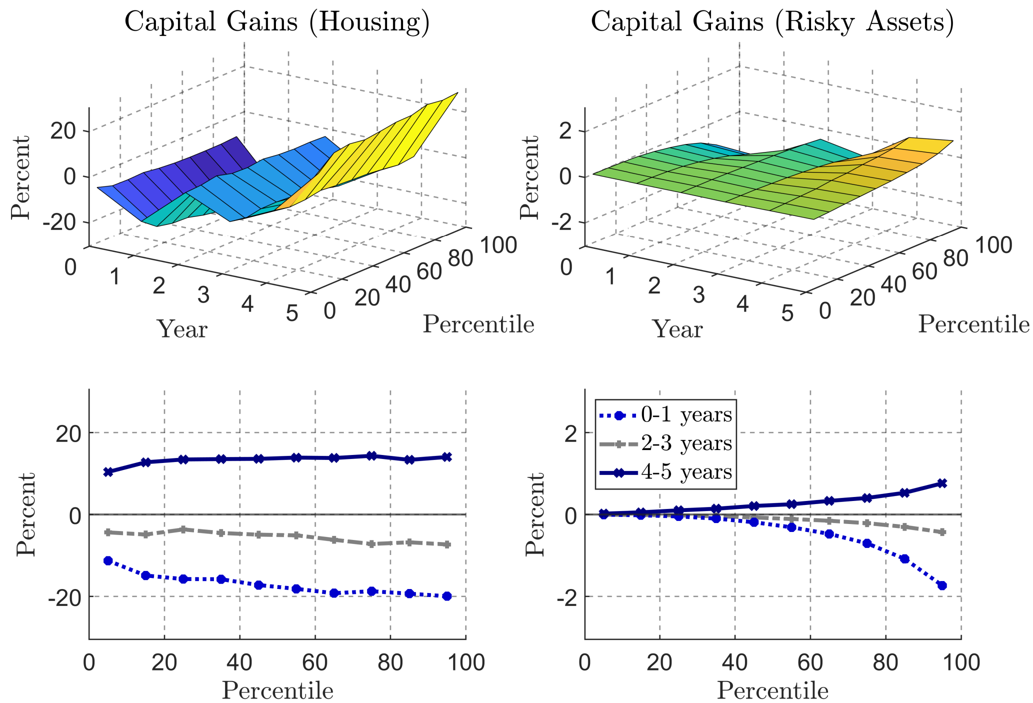

Figure 14: **Impulse Responses of Capital Gains in Housing and Risky Financial Assets.**
Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures C.2-C.5 in Appendix C.2 for group responses and group comparisons with confidence bands.

Holm and Natvik (2021) as shown in Figure C.7. Even households at the top of the liquid asset distribution may be affected by borrowing constraints if their income and consumption is large compared to their liquidity holdings. To exclude this possibility, Figure C.8 demonstrates that the results are nearly unchanged if we order households by lagged liquid asset holdings per unit of income. To verify that our procedure of imputing capital gains does not drive our results, we consider two exercises. First, we exclude equity holders from the estimation in Figure C.9. Second, we consider consumption responses based on (extreme) assumptions about capital gains in stocks in Figure C.10.[33] In both cases, our results are robust. Finally, Figures C.11 and C.12 show that the responses described above are not merely an artifact of age effects or initial income differences that are correlated with liquid asset holdings. Splitting the liquidity deciles additionally into quintiles of the age or income distribution yields robust estimates across the subgroups.

An alternative explanation for the behavior at the top of the distribution could be a reverse causality between liquid asset holdings and large consumption expenditures. Households may hold liquid assets, because they plan to make large purchases, for example, of durables in the future. The strong positive consumption response of the top 10 percent in the first years may then reflect

[33]Note that we observe housing transactions and therefore the active component of saving in housing wealth. Hence, any potential mismeasurement in the capital gains on housing units does not affect our consumption measure. For stocks, our method of imputing capital gains ensures that the sum of individual capital gains adds up to aggregate capital gains in the stock market (before 2005) or an individual stock (after 2005). While aggregate capital gains are matched by construction, a concern could be that we distribute capital gains incorrectly across asset holders. Appendix C.3.2 demonstrates that measurement error in the capital gains on stocks does not lie at the heart of our findings.

that the additional interest income received induces them to bring the durable purchase forward. However, for at least two reasons this channel is unlikely to play a significant role. First, the amount of liquid assets required to be in the top 10 percent is substantial (median of about 100,000 U.S. dollars in our sample). Second, a majority of households in the top group remain among the top 10 percent every year.[34] The high persistence in liquid asset holdings makes it unlikely that households with large planned expenditures are an important source of the consumption response that we observe among high-liquidity households.

The magnitude of the responses to monetary policy shocks in units of lagged household income can differ across the liquid asset distribution due to systematic differences in both the change in the variable of interest and the level of household income prior to the shock. In Appendix C.4, we reestimate equation (5) replacing $inc_{i,t-1}$ with average disposable income across all households to investigate the influence of the normalization. Since this makes the impulse responses summable across liquid asset groups, it also allows us to gauge whether household groups exist that are of disproportionate importance for the aggregate consumption response. The results confirm the significance of the households at the top of the liquid asset distribution, which show a large response relative to their initial income, as seen before, and in addition have comparably high income on average.

4.3 Cash-Flow Effects

To carve out the effect of interest rate-sensitive cash flows in more detail, we reorder households according to their net interest rate exposure and form groups based on the position of households in the resulting distribution before reestimating the local projections in (5). Net interest rate exposure is measured as deposits net of debt. Recall from Section 3 that most debt is held in the form of adjustable-rate mortgages and that both deposit and loan rates are highly elastic to the policy rate. Our measure of interest rate exposure is closely related to the concept of "unhedged interest rate exposure" in Auclert (2019). Note, however, that it abstracts from the maturity and duration of assets and liabilities.

The results are shown in Figures 15 and 16 with net debtors at the bottom of the distribution and net creditors at the top. The full impulse responses together with statistical significance tests can be found in Appendix C.5. In the year in which the shock occurs and in the year after that, the consumption responses at the top and at the bottom of the distribution are of similar magnitude but opposite sign. Both are significant at the 5 percent level in year 1. Households with high positive net interest rate exposure tend to hold large amounts of interest rate-sensitive assets but little or no debt. Their capital revenues increase sharply and highly significantly in response to a positive interest rate shock while their capital expenditures remain unaffected. The strong reaction of net financial income overcompensates a small but significant decline in nonfinancial income, implying that disposable income rises at the top of the distribution. At the other end of the spectrum,

[34]On average, about 67 percent of the households in the top group in a given year were also among the top 10 percent in all of the previous five years.

households with high negative interest rate exposure frequently hold little deposits and bonds but have large debt positions. Their net financial income response is nearly the mirror image of the one by households at the top of the distribution. The large decline in net financial income resulting from a strong and statistically significant increase in capital expenditures overcompensates a small rise in nonfinancial income. Disposable income therefore declines. At the top of the distribution, about half of the income change is saved in years 0 and 1. In contrast, consumption moves nearly one-for-one with disposable income at the bottom.

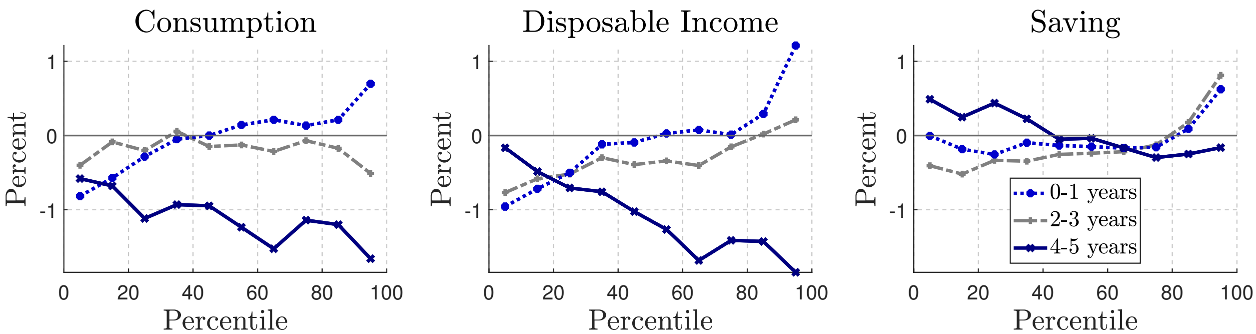

Figure 15: **Impulse Responses by Net Interest Rate Exposure.**
Notes: Changes relative to lagged net income at segments of the distribution of net interest rate exposure. See Figures C.23-C.26 in Appendix C.5 for group responses and group comparisons with confidence bands.

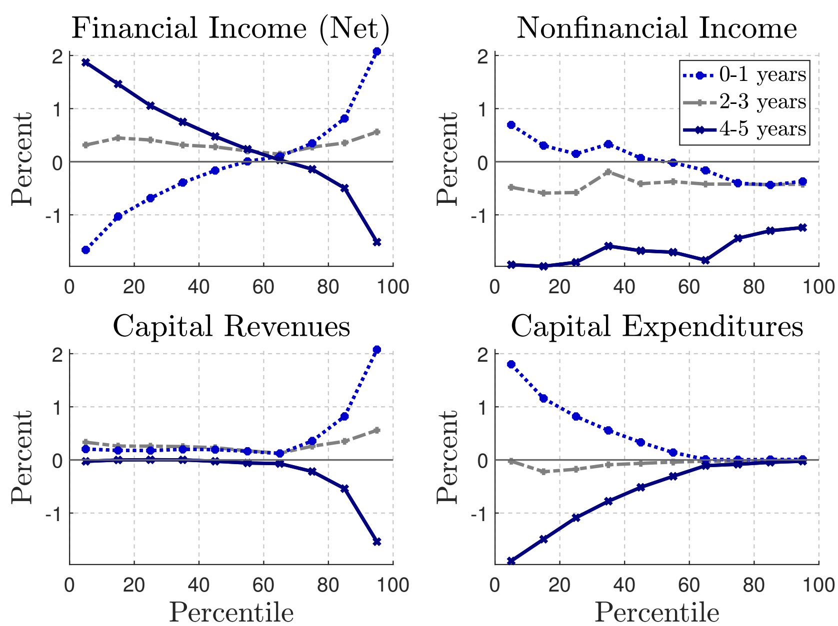

Figure 16: **Impulse Responses of Income Components by Net Interest Rate Exposure.**
Notes: Changes relative to lagged net income at segments of the distribution of net interest rate exposure. See Figures C.23-C.26 in Appendix C.5 for group responses and group comparisons with confidence bands.

The symmetry between the responses of households with large positive and large negative interest rate exposure disappears over time. While the financial income responses reverse after four to five years as the policy rate dips below its initial level, the responses of nonfinancial income are significantly negative across the entire distribution, as above. At the bottom, the sharp decline in capital expenditures is nearly offset by the fall in nonfinancial income. Since saving increases somewhat, the consumption response after four to five years is negative but small for households at the bottom. The increase in saving could be explained by debt prepayment in times of low debt-servicing costs. At the top, the financial and nonfinancial income responses are significantly negative in both years, implying that disposable income strongly declines. As the saving response is muted, consumption falls substantially more than at the bottom of the distribution.

The households can be reordered along many other dimensions. Doing so implies that net interest rate exposure is generally more dispersed within groups but less across groups. Accordingly, there are less clear patterns in the transmission of monetary policy to financial income and therefore to consumption. Figure 17 illustrates this observation for the income and the net worth distribution.

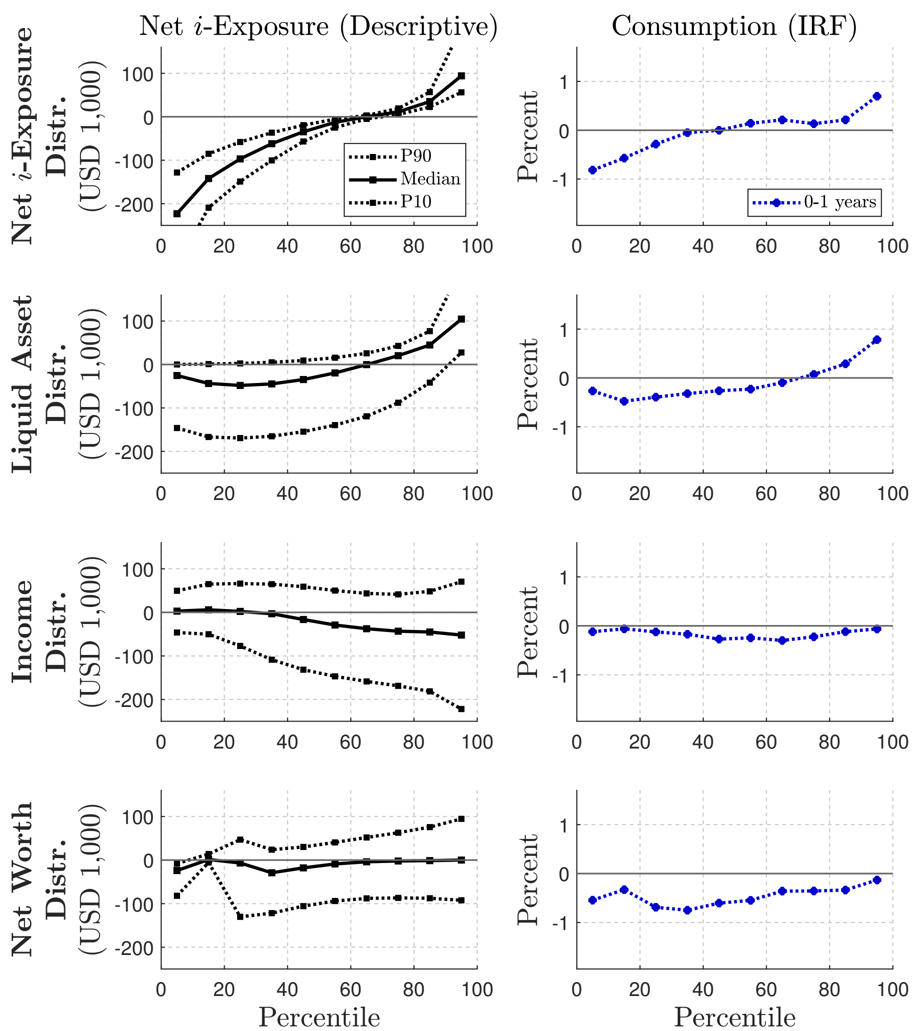

Figure 17: **Net Interest Rate Exposure and Consumption Responses for Various Distributions.**
Notes: Median, 10th percentile, and 90th percentile of net interest rate exposure by group along the distributions of net interest rate exposure, liquid assets, income, and net worth (left column). Impulse responses in changes relative to lagged net income at segments of the same distributions (right column).

The panels in the left column show descriptive statistics of net interest rate exposure in each of the ten groups when households are ordered by their net interest rate exposure, liquid assets, income, and net worth in the same way as above. The panels in the right column contain the corresponding consumption responses in year 0 to 1. All income groups include households with substantial positive and negative exposure as can be seen from the dashed lines indicating the 10th and the 90th percentile. The large disparities within the income groups imply that the median exposure differs less strongly across groups than when households are ordered by their net interest rate exposure or their liquid asset holdings. The profile of consumption responses across the income distribution is therefore nearly flat. An analogous point applies to the net worth distribution demonstrating the importance of distinguishing between portfolio positions with interest rate sensitive and insensitive cash flows.

Our results are in line with the findings of other studies that evaluate the impact of cash flow adjustments induced by interest rate changes on household consumption. Di Maggio et al. (2017) examine the effects of the large declines in interest payments experienced by a number of households in the United States with adjustable-rate mortgages in the aftermath of the Great Recession. They show that households that were affected responded with sizable increases in durable consumption, namely car purchases, and voluntary repayment of mortgage debt. Based on data from Sweden, Floden et al. (2020) equally find evidence for interest rate changes to feed into household consumption through adjustable-rate mortgages.

5 Direct and Indirect Effects of Monetary Policy

The sizable nonfinancial income responses uncovered before and their close relation with households' consumption responses suggest that indirect effects are important for the transmission of monetary policy, as argued by Kaplan, Moll and Violante (2018). In this section, we analyze whether that is the case by separating households' consumption responses into direct and indirect effects.

To distinguish direct and indirect effects empirically, we estimate two separate types of consumption responses to monetary policy shocks. The first estimate is the one based on the regression setup (4) in Section 4.1, which considers all households and the associated responses include both direct and indirect effects. The second estimate is based on the same specification but also controls for a household's income change over the respective impulse response horizon. In particular, we consider local projections of the form

$$\frac{c_{i,t+h} - c_{i,t-1}}{\bar{c}_{t-1}} = \delta_i^h + \beta^h \epsilon_t^{MP} + \sum_{m=1}^{3} \mu_m^h X_{i,t-m} + \sum_{k=1}^{K} \sum_{m=0}^{h} \gamma_m^{h,k} \tilde{y}_{i,t+m}^k \quad , \tag{6}$$

for $h = 0, 1, ..., 5$. The dependent variable is given by the change in consumption for household i from period $t-1$ to $t+h$, scaled by the average consumption level in $t-1$ across the whole sample. As above, we include a household fixed effect δ_i^h and the same controls $X_{i,t}$ as in (4), that is, three years of lagged values of the monetary policy shock and two years of lagged one-year growth rates of the dependent variable.

The only difference between (4) and (6) is the additional set of income controls $\sum_{k=1}^{K}\sum_{m=0}^{h}\gamma_m^{h,k}\widetilde{y}_{i,t+m}^k$, where $\widetilde{y}_{i,t+m}^k = \left(y_{i,t+m}^k - y_{i,t-1}^k\right)/\bar{c}_{t-1}$ is a particular subcomponent k of disposable income.[35] The separate inclusion of various income types allows for the possibility that the income elasticities $\gamma_m^{h,k}$ vary by income type. Importantly, the timing of the income controls differs from the timing of the remaining regressors. For the impulse response horizon h, specification (6) controls for changes in household income that occur between $t-1$ and $t+h$. To allow for flexible dynamic relations between income and consumption changes, we decompose the income changes into $(h+1)$ changes for the response up to horizon h.

The estimated coefficients β^h in (6) therefore give the partial effect of monetary policy on consumption at horizon h holding income constant over the same time period. We exclude interest income and interest expenses from the set of income controls since changes in these items are considered direct income effects in Kaplan, Moll and Violante (2018). Hence, β^h in (6) gives an estimate of the direct effects of monetary policy. In turn, any differences between the estimated coefficients β^h in (4) and those in (6) are associated with indirect effects of monetary policy.

Our decomposition closely follows the theoretical separation by Kaplan, Moll and Violante (2018) (see, e.g., equation (3) in their paper). In Appendix E.1, we show that our empirical decomposition can be interpreted as an omitted variable bias problem, arising from the omission of the income controls in (4). To identify the direct effect using (6), two conditions have to be satisfied. First, households' consumption changes cannot feed back into their income within the same period, otherwise (6) would be subject to a reverse causality problem. For example, by lowering their consumption and increasing their saving, households could raise their interest income and may consume out of the additional income within the same period. However, given the exclusion of interest income and expenses from the set of income controls, such reverse causality likely plays a minor role. Second, to achieve identification, household income needs to move due to structural variation that does not also affect household consumption directly, unless it enters as a control in (6). Idiosyncratic income shocks, which are generally incorporated into HANK-models, are a prime example of such variation. Below, we show that our baseline results are robust to using lottery winnings as an instrument for idiosyncratic income movements.

While our decomposition closely resembles the one in Kaplan, Moll and Violante (2018), there are also three important differences. First, the indirect effect in Kaplan, Moll and Violante (2018) includes changes in consumption that are due to anticipated future changes in income. In contrast, we control only for changes in income up to horizon h in (6), since the data only cover realized income at the household level and not household expectations of future income paths. Second,

[35] All income changes are scaled by the average consumption level $\bar{c}_{t-1}$ in period $t-1$ to match the dependent variable. In various robustness checks, we confirmed that this scaling of the controls does not impact the results.

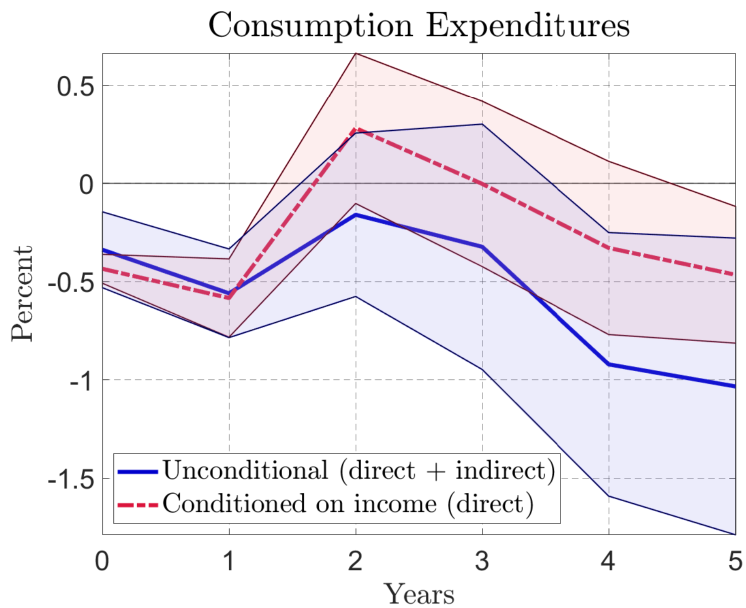

Figure 18: **Direct and Indirect Effects of Monetary Policy.**
Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at an annual frequency, based on the local projection approaches in (4) and (6). The blue line shows the estimated impulse responses without controlling for income, the red dashed line shows the responses with income controls. 68 percent confidence bands shown, using Driscoll and Kraay (1998) standard errors.

the indirect effect in Kaplan, Moll and Violante (2018) also includes indirect wealth effects. In a robustness exercise in Appendix E, we therefore control for changes of household wealth from risky assets. The extra capital gains controls do not affect our results. Third, Kaplan, Moll and Violante (2018) decompose the contemporaneous and first-year consumption responses into direct and indirect effects. In contrast, we analyze the path of consumption up to five years after the shock. This difference is important since the aggregate consumption response in the data builds up over time (see Figure 7 and, e.g., Christiano, Eichenbaum and Evans, 2005).[36]

The results are shown in Figure 18.[37] The blue line shows the estimated impulse response of consumption without the income controls and the red dashed line shows those with income controls. For the contemporaneous and the subsequent year, the two lines lie nearly on top of each other. Hence, any consumption change in response to a monetary policy shock can be associated with direct effects. That is, the rise in the policy rate over this period (as shown in Figure 3) leads to an increase in saving and a fall in consumption that are unrelated to a change in nonfinancial income over this period. After year two, the two lines start to diverge. The red dashed line even turns positive in year two after a shock, which can be associated with the fall in the policy rate around this time (see Figure 3). At the five-year horizon, indirect effects dominate, accounting for more

[36] That is also a difference in comparison with the decomposition by Auclert (2019). He considers one-period changes in income, prices, and interest rates to separate the aggregate contemporaneous consumption response into different channels.

[37] For these estimations, only households that are observed for all periods over the five-year impulse response horizon are included in the sample to ensure that the estimated differences are not due to different samples.

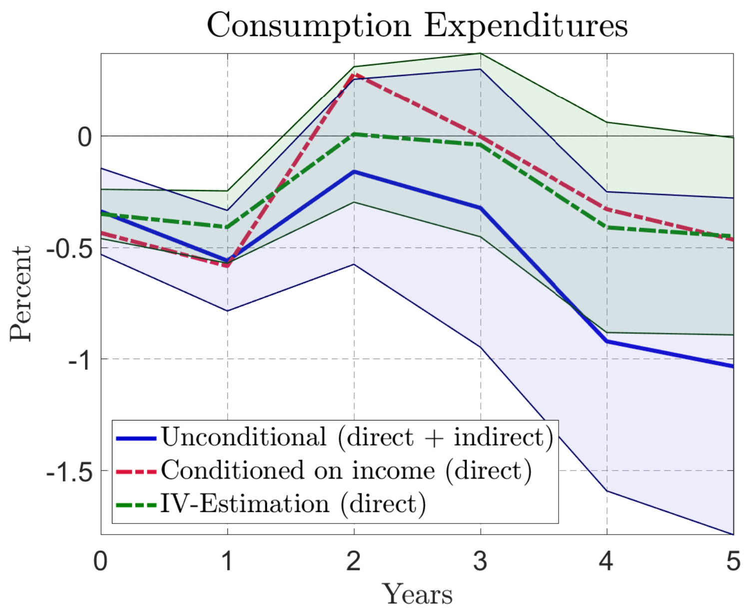

Figure 19: **Direct and Indirect Effects of Monetary Policy − IV-Estimation.**
Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at an annual frequency based on the local projection approach in (4) (blue solid line), the instrumental variable approach described in Appendix E.1 (green dashed line), and the local projection approach in (6) (red dashed line). 68 percent confidence bands shown, using Driscoll and Kraay (1998) standard errors.

than 50 percent of the consumption response.[38]

Thus, our findings show that the dominance of the indirect effect of monetary policy, described theoretically by Kaplan, Moll and Violante (2018), also holds up empirically. However, in Kaplan, Moll and Violante (2018), the indirect effect already comes into play over the first year of the shock. In contrast, our results show that it takes time for the indirect effect to build up. Initially, the consumption response is almost completely explained by direct effects. After around two to three years, the indirect effect dominates and explains most of the consumption response in the data. In Figure E.2 in Appendix E.4, we repeat the above exercise across the liquid asset distribution. Direct effects dominate initially for all households across the entire distribution. After two to three years, indirect effects account for the majority of the consumption response for the bottom half of the distribution, for which nonfinancial income is particularly important. In fact, for the bottom 10 percent, indirect effects explain nearly all of the drop in consumption.[39]

A potential identification concern is that household income may be driven by structural variation that also affects consumption directly and that is not controlled for in (6). If such variation

[38] In comparison, Kaplan, Moll and Violante (2018) show that indirect effects account for around 80 percent of the total consumption response over the first year in their model.

[39] Based on these regressions, Figure E.4 in Appendix E.5 shows the estimated coefficient with respect to the contemporaneous nonfinancial income control for the time zero impulse response horizon $\hat{\gamma}_0^{0,k}$ based on regression (6) across the liquid asset distribution. The estimates range from around 0.7 at the bottom to around 0.3 at the top of the distribution. While these coefficients should not be understood as MPC measures, they line up closely with the MPC estimates that are shown in Figure 10.

is large, regression (6) does not identify the true direct effect. To address this concern, we use lottery winnings as instruments for household idiosyncratic income shocks based on the data in Fagereng, Holm and Natvik (2021). We describe the data and the procedure in Appendix E.2. In the first stage, household nonfinancial income is projected on the monetary policy shock and the lottery winnings, in addition to our standard controls in (6).[40] In the second stage, the predicted values are used as controls as in (6). Figure 19 shows the results. Strikingly, the estimated impulse responses in Figures 18 and 19 are nearly the same.

6 Conclusion

The frequently estimated aggregate responses to monetary policy mask large heterogeneity at the household level. Our results about the effects of monetary policy on household consumption, income, and saving based on administrative data from Norway may guide policy makers and a new generation of HANK models alike.

We order households according to their liquid asset positions and estimate impulse responses at fine segments along the resulting distribution. Households at the bottom of the distribution show a stronger consumption response than those around the median. Following a monetary tightening, low-liquidity households reduce their consumption when their disposable income starts to fall, while households around the median initially save less or borrow more. These findings are consistent with obstacles to consumption smoothing, such as a borrowing constraint, faced by households with low liquid asset holdings, a mechanism that lies at the heart of many recent HANK models.

Monetary policy affects the net financial income of borrowers and savers. Our estimates indicate that these income changes are permitted to feed into consumption. When the policy rate rises, net creditors see their interest income increase, whereas net borrowers experience a rise in their interest expenses. Both of these effects have a substantial impact on household consumption, giving rise to an initial response that is of opposite sign at both ends of the liquid asset distribution. The finding that high-liquidity households increase their consumption in response to the rise in interest income associated with a monetary tightening is perhaps most surprising. It presents a challenge for HANK models: even households with substantial liquidity consume a sizable fraction of temporary income changes.

We expect the heterogeneity in responses that we uncover to be a feature of monetary policy in many countries beyond Norway. The distribution of assets and debt is more unequal in countries like the United States than in Norway, which by itself points to even stronger heterogeneity in the effects on consumption. The prevalence of adjustable-rate mortgages and the relatively high elasticity of deposit rates to monetary policy suggest that the household cash-flow channel may

[40]The F-test statistics from the various first-stage regressions are all well above typical critical values testing for weak identification, because the lottery winnings directly enter household nonfinancial income.

be somewhat more pronounced in Norway than in countries where mortgages predominantly carry a fixed rate and interest rate pass-through is more subdued. However, since HANK models typically abstract from frictional interest rate pass-through, Norway is a suitable environment for testing the predictions of these models.

The relative strength of the direct and indirect effects of monetary policy depends on the impulse response horizon. Direct effects dominate in the first two years after a shock. Indirect effects gain in importance in the following years as the response of nonfinancial income builds up, suggesting that the characteristics of a country's productive sector and labor market affect monetary policy transmission to private households. Our results confirm that indirect effects play a sizable role in transmitting interest rate changes to consumption, with the important qualification that it takes time for these effects to unfold.

References

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109(6):2333–67.

Auclert, Adrien, Matthew Rognlie and Ludwig Straub. 2018. The Intertemporal Keynesian Cross. Unpublished manuscript.

Auclert, Adrien, Matthew Rognlie and Ludwig Straub. 2020. Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model. Working Paper 26647 National Bureau of Economic Research.

Baker, Scott R, Lorenz Kueng, Steffen Meyer and Michaela Pagel. 2018. Measurement error in imputed consumption. Unpublished manuscript.

Barnichon, Regis and Christian Matthes. 2018. "Functional Approximation of Impulse Responses." *Journal of Monetary Economics* 99(C):41–55.

Bayer, Christian, Ralph Luetticke, Lien Pham-Dao and Volker Tjaden. 2019. "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk." *Econometrica* 87(1):255–290.

Beraja, Martin, Andreas Fuster, Erik Hurst and Joseph Vavra. 2018. "Regional Heterogeneity and the Refinancing Channel of Monetary Policy." *Quarterly Journal of Economics* 134(1):109–183.

Berger, David, Konstantin Milbradt, Fabrice Tourre and Joseph Vavra. 2018. Mortgage Prepayment and Path-Dependent Effects of Monetary Policy. Working Paper 25157 National Bureau of Economic Research.

Bilbiie, Florin O. 2008. "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic." *Journal of Economic Theory* 140(1):162–196.

Bilbiie, Florin O. 2020. "The new Keynesian cross." *Journal of Monetary Economics* 114:90–108.

Bilbiie, Florin O., Diego R. Känzig and Paolo Surico. 2020. Capital, Income Inequality, and Consumption: The Missing Link. Unpublished manuscript.

Bjørnland, Hilde C, Thomas Ekeli, Petra M Geraats and Kai Leitemo. 2004. "Norges Bank Watch 2004, An Independent Review of Monetary Policymaking in Norway." *Norges Bank Watch Report Series No 5* .

Bunn, Philip, Jeanne Le Roux, Kate Reinold and Paolo Surico. 2018. "The consumption response to positive and negative income shocks." *Journal of Monetary Economics* 96:1–15.

Campbell, John Y. and N. Gregory Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." *NBER Macroeconomics Annual* 4:185–216.

Challe, Edouard. 2020. "Uninsured unemployment risk and optimal monetary policy in a zero-liquidity economy." *American Economic Journal: Macroeconomics* 12(2):241–83.

Chodorow-Reich, Gabriel, Plamen T Nenov and Alp Simsek. 2021. "Stock Market Wealth and the Real Economy: A Local Labor Market Approach." *forthcoming, American Economic Review* .

Christiano, Lawrence J., Martin Eichenbaum and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1):1–45.

Cloyne, James, Clodomiro Ferreira and Paolo Surico. 2020. "Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism." *Review of Economic Studies* 87(1):102–129.

Cloyne, James and Patrick Huertgen. 2016. "The Macroeconomic Effects of Monetary Policy: A New Measure for the United Kingdom." *American Economic Journal: Macroeconomics* 8(4):75–102.

Coibion, Olivier. 2012. "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Macroeconomics* 4(2):1–32.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng and John Silvia. 2017. "Innocent Bystanders? Monetary Policy and Inequality in the U.S." *Journal of Monetary Economics* 88:70–89.

Crawley, Edmund and Andreas Kuchler. 2020. Consumption Heterogeneity: Micro Drivers and Macro Implications. Finance and Economics Discussion Series 2020-005 Board of Governors of the Federal Reserve System.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita and Michael Weber. 2019. Human Frictions to the Transmission of Economic Policy. 2019 Meeting Papers 339 Society for Economic Dynamics.

Debortoli, Davide and Jordi Gali. 2018. Monetary Policy with Heterogeneous Agents: Insights from TANK models. Unpublished manuscript.

Den Haan, Wouter J, Pontus Rendahl and Markus Riegler. 2017. "Unemployment (Fears) and Deflationary Spirals." *Journal of the European Economic Association* 16(5):1281–1349.

Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru and Vincent Yao. 2017. "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging." *American Economic Review* 107(11):3550–88.

Di Maggio, Marco, Amir Kermani and Kaveh Majlesi. 2020. "Stock market returns and consumption." *The Journal of Finance* 75(6):3175–3219.

Drechsler, Itamar, Alexi Savov and Philipp Schnabl. 2017. "The Deposits Channel of Monetary Policy." *Quarterly Journal of Economics* 132(4):1819–1876.

Driscoll, John C. and Aart C. Kraay. 1998. "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data." *Review of Economics and Statistics* 80(4):549–560.

Eichenbaum, Martin, Sergio Rebelo and Arlene Wong. 2019. State Dependent Effects of Monetary Policy: The Refinancing Channel. Unpublished manuscript.

Eika, Lasse, Magne Mogstad and Ola L Vestad. 2020. "What can we learn about household consumption expenditure from data on income and assets?" *Journal of Public Economics* 189:104163.

Fagereng, Andreas and Elin Halvorsen. 2017. "Imputing consumption from Norwegian income and wealth registry data." *Journal of Economic and Social Measurement* 42(1):67–100.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino and Luigi Pistaferri. 2020. "Heterogeneity and persistence in returns to wealth." *Econometrica* 88(1):115–170.

Fagereng, Andreas, Martin B. Holm, Benjamin Moll and Gisle Natvik. 2019. Saving Behavior across the Wealth Distribution: The Importance of Capital Gains. Working Paper 26588 National Bureau of Economic Research.

Fagereng, Andreas, Martin B. Holm and Gisle Natvik. 2021. "MPC Heterogeneity and Household Balance Sheets." *Forthcoming, American Economic Journal: Macroeconomics* .

Fagereng, Andreas, Martin B Holm and Kjersti N Torstensen. 2020. "Housing Wealth in Norway, 1993 - 2015." *Journal of Economic and Social Measurement* 45(1):65–81.

Floden, Martin, Matilda Kilstrom, Josef Sigurdsson and Roine Vestman. 2020. "Household Debt and Monetary Policy: Revealing the Cash-Flow Channel." *forthcoming, The Economic Journal* .

Gabaix, Xavier and David Laibson. 2001. The 6D Bias and the Equity-Premium Puzzle. In *NBER Macroeconomics Annual 2001, Volume 16*. NBER Chapters National Bureau of Economic Research, Inc pp. 257–330.

Galí, Jordi, David López-Salido and Javier Vallés. 2007. "Understanding the Effects of Government Spending on Consumption." *Journal of the European Economic Association* 5(1):227–270.

Gertler, Mark and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7(1):44–76.

Gornemann, Nils, Keith Kuester and Makoto Nakajima. 2016. Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy. Unpublished manuscript.

Guerkaynak, Refet, Brian Sack and Eric Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking* 1(1):55–93.

Guerrieri, Veronica and Guido Lorenzoni. 2017. "Credit Crises, Precautionary Savings, and the Liquidity Trap." *Quarterly Journal of Economics* 132(3):1427–1467.

Hagedorn, Marcus, Jinfeng Luo, Iourii Manovskii and Kurt Mitman. 2019. "Forward guidance." *Journal of Monetary Economics* 102(C):1–23.

Hansen, Bruce E. 2019. *Econometrics*. University of Wisconsin.

Jappelli, Tullio and Luigi Pistaferri. 2010. "The Consumption Response to Income Changes." *Annual Review of Economics* 2(1):479–506, 09.

Johnson, David S., Jonathan A. Parker and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96(5):1589–1610.

Jordà, Oscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95(1):161–182.

Kaplan, Greg, Benjamin Moll and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108(3):697–743.

Kaplan, Greg and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82(4):1199–1239.

Kaplan, Greg, Giovanni L. Violante and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity* 45(1):77–153.

Kekre, Rohan and Moritz Lenel. 2020. Monetary Policy, Redistribution, and Risk Premia. Unpublished manuscript.

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

Leth-Petersen, Søren. 2010. "Intertemporal consumption and credit constraints: Does total expenditure respond to an exogenous shock to credit?" *American Economic Review* 100(3):1080–1103.

Lintner, John. 1956. "Distribution of Incomes of Corporations Among Dividends, Retained Earnings, and Taxes." *American Economic Review* 46(2):97–113.

Luetticke, Ralph. 2020. "Transmission of Monetary Policy with Heterogeneity in Household Portfolios." *forthcoming, American Economic Journal: Macroeconomics* .

Mankiw, N. Gregory. 2000. "The Savers-Spenders Theory of Fiscal Policy." *American Economic Review* 90(2):120–125.

Mankiw, N. Gregory and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117(4):1295–1328.

McKay, Alisdair, Emi Nakamura and Jon Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review* 106(10):3133–58.

McKay, Alisdair and Ricardo Reis. 2016. "The Role of Automatic Stabilizers in the U.S. Business Cycle." *Econometrica* 84(1):141–194.

Mertens, Karel and Morten O. Ravn. 2013. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *American Economic Review* 103(4):1212–1247.

Misra, Kanishka and Paolo Surico. 2014. "Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs." *American Economic Journal: Macroeconomics* 6(4):84–106.

Newey, Whitney and Kenneth West. 1987. "A Simple, Positive-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55:703–708.

Oh, Hyunseung and Ricardo Reis. 2012. "Targeted transfers and the fiscal response to the great recession." *Journal of Monetary Economics* 59(S):50–64.

Olivei, Giovanni and Silvana Tenreyro. 2007. "The Timing of Monetary Policy Shocks." *American Economic Review* 97(3):636–663.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103(6):2530–53.

Patterson, Christina. 2018. The Matching Multiplier and the Amplification of Recessions. Unpublished manuscript.

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

Ravn, Morten O. and Vincent Sterk. 2017. "Job uncertainty and deep recessions." *Journal of Monetary Economics* 90(C):125–141.

Romer, Christina D. and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94(4):1055–1084.

Slacalek, Jiri, Oreste Tristani and Giovanni L. Violante. 2020. "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.

Stock, James H. and Mark W. Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics." *Economic Journal* 128(610):917–948.

Wong, Arlene. 2019. Refinancing and the Transmission of Monetary Policy to Consumption. Unpublished manuscript.