

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

The Transmission of Monetary Policy under the Microscope

Martin Blomhoff Holm
University of Oslo

Pascal Paul
Federal Reserve Bank of San Francisco

Andreas Tischbirek
University of Lausanne

April 2020

Working Paper 2020-03

<https://www.frbsf.org/economic-research/publications/working-papers/2020/03/>

Suggested citation:

Holm, Martin Blomhoff, Pascal Paul, Andreas Tischbirek. 2020. “The Transmission of Monetary Policy under the Microscope,” Federal Reserve Bank of San Francisco Working Paper 2020-03. <https://doi.org/10.24148/wp2020-03>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

The Transmission of Monetary Policy under the Microscope

Martin Blomhoff Holm, Pascal Paul, Andreas Tischbirek*

April 2020

Abstract

We investigate the transmission of monetary policy to household consumption using detailed administrative data on the universe of households in Norway. Based on a novel series of identified monetary policy shocks, we estimate the dynamic responses of consumption, income, and saving along the liquid asset distribution of households. We find that low-liquidity but also high-liquidity households show strong responses, interest rate changes faced by borrowers and savers feed into consumption, and indirect effects of monetary policy outweigh direct effects, albeit with a delay. Overall, the results support the importance of financial frictions, cash-flow channels, and heterogeneous effects of monetary policy.

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

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

*We thank Florin Bilbiie, James Cloyne, Andreas Fagereng, Kilian Huber, Òscar Jordà, Greg Kaplan, Ben Moll, Chris Moser, Emi Nakamura, Christina Patterson, Matthew Rognlie, and Jón Steinsson for detailed comments and seminar participants at UC Berkeley, the Board of Governors of the Federal Reserve, the University of Chicago, Columbia Business School, the Federal Reserve Bank of San Francisco, the University of Lausanne, Norges Bank, and the University of Oslo 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. The article was written as part of a research project at Statistics Norway supported by the Research Council of Norway, project number 287720. We 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.

Martin Blomhoff Holm: University of Oslo, Email: martin.b.holm@outlook.com.

Pascal Paul: Federal Reserve Bank of San Francisco, Email: pascal.paul@sf.frb.org.

Andreas Tischbirek: HEC Lausanne, University of Lausanne, Email: andreas.tischbirek@unil.ch.

1 Introduction

In recent years, monetary economics has experienced a profound change. Traditionally, the large majority of models used to analyze monetary policy assumed away all heterogeneity among households. Empirical studies equally focused on the impact of monetary policy on macroeconomic aggregates. A new class of models has recently emerged that recognizes the sizable observed heterogeneity among households and its significance for the transmission of monetary policy. Most prominently, the Heterogeneous Agent New Keynesian (HANK) model by [Kaplan, Moll and Violante \(2018\)](#) allows for rich heterogeneity in household characteristics such as consumption, income, and wealth.

Recent 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, 2019](#)). 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.¹ 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 successful at generating MPCs at the aggregate level that are closer to the values estimated from empirical data. In the presence of borrowing costs and constraints, income changes of households with few liquid assets are imperfectly smoothed and permitted to feed into consumption.² [Kaplan and Violante \(2014\)](#) and [Kaplan, Moll and Violante \(2018\)](#) 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 model that allows for a substantial fraction of households with few liquid assets, [Kaplan, Moll and Violante \(2018\)](#) show 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 novel view of the monetary transmission mechanism implied by HANK models is scarce. 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 sheets, income, and consumption of households. We therefore turn to a country that collects

¹See, 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 \(2018\)](#).

²The important role of liquidity in the presence of financial frictions is not specific to the model by [Kaplan, Moll and Violante \(2018\)](#), but lies at the heart of many models in the HANK literature. See [Kuester, Gornemann and Nakajima \(2016\)](#), [McKay, Nakamura and Steinsson \(2016\)](#), [McKay and Reis \(2016\)](#), [Guerrieri and Lorenzoni \(2017\)](#), [Debortoli and Gali \(2018\)](#), [Ravn and Sterk \(2018\)](#), [Bayer, Luetticke, Pham-Dao and Tjaden \(2019\)](#), [Bilbiie \(2019\)](#), and [Luetticke \(2019\)](#) among others.

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

rich information from its inhabitants: Norway. Specifically, we draw on administrative data that contain records of the income and wealth of the entire universe of households residing 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 unique 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 novel 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. Consumption responses move closely with 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 findings are not consistent with the predictions of standard HANK models. In these models, the response of high-liquidity households to interest rate changes is mainly determined by intertemporal substitution. In contrast, we find that households with large liquidity positions increase both consumption and saving after a monetary tightening, before consumption ultimately falls. We show that the sizable increases on impact are related to a rise in financial income, in particular interest income, which is directly affected by the policy rate. Since temporary changes in disposable income substantially feed into consumption, our estimates indicate that households at the top of the liquid asset distribution have sizable MPCs and that their responses are not primarily driven by an intertemporal substitution motive.

To gauge if specific groups of households drive the responses to monetary policy at the aggregate level, we ask how much of the aggregate consumption response is explained by households with

different levels of liquid asset holdings. While households at the bottom of the liquid asset distribution show an elevated response relative to their own consumption level, they only account for a small share of total consumption in the economy. Correspondingly, their contribution to the aggregate response is close to that of the median household. More broadly, below the 80th percentile, all deciles of the distribution contribute a roughly equal share to the total response. The top 20 percent of households, by contrast, make a significantly bigger contribution since they account for a large share of aggregate consumption. By breaking down the top 10 percent further, we show that this pattern is not driven by a handful of households with exorbitant liquid wealth.

Interest income and expenditures 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 by households at the bottom (net borrowers) is the mirror-image of that by households at the top (net creditors). The strong effect at both ends is explained by the prevalence of adjustable-rate mortgages and the high flexibility of deposit rates in Norway. Cash-flow adjustments feed into household consumption, which can be attenuated or magnified by the response of nonfinancial income.

Our estimates 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 provide the first comprehensive empirical estimates of the relative importance of direct and indirect effects. The relative size of the two types of 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 \(2019\)](#) 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. \(2017\)](#), 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. [Ampudia et al. \(2018\)](#) and [Slacalek, Tristani and Violante \(2020\)](#) use variations of [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 various 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 such monetary policy shocks.⁴ The idea of this approach is to orthogonalize 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 21](#) 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

$$\begin{aligned} \Delta i_m = & \alpha_1 + \alpha_2 i_{m,-1} + \sum_{k=0}^1 \beta_k^\pi \pi_{t+k}^m + \sum_{k=0}^1 \beta_k^{\Delta\pi} \Delta \pi_{t+k}^m \\ & + \sum_{k=0}^1 \beta_k^y y_{t+k}^m + \sum_{k=0}^1 \beta_k^{\Delta y} \Delta y_{t+k}^m + \gamma_1 ex_{m,-1} + \gamma_2 I_{IT} \cdot ex_{m,-1} + \epsilon_m^{MP} \quad , \end{aligned} \quad (1)$$

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 within period t .⁵ Following [Romer and Romer \(2004\)](#), we include forecasts for GDP y_{t+k}^m and the CPI π_{t+k}^m for horizon $t+k$ and the corresponding forecast changes, denoted $\Delta \pi_{t+k}^m$ and Δy_{t+k}^m . More details about these forecasts and their timing relative to the policy meetings are provided in the next section.

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 the 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 this change, the central bank to some extent 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_{IT}

⁴A 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, 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.

⁵Occasionally, 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.

⁶The 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.

that takes the value of one for the pre-inflation targeting era.⁷ The residual ϵ_m^{MP} is understood as a measure of the monetary policy shock associated with meeting m . In Section 2.4, we check (and confirm) the robustness of the estimated shocks and impulse responses to a number of variations of (1).

2.1 Forecasts

Besides information on all policy meetings and the associated rate changes between 1994 and 2018, we also collect data on the forecasts constructed by Norges Bank. From 1996:M11 until 2018:M12, the forecasts are obtained from reports that are publicly available. Until 2006, Norges Bank published its forecasts in the so-called Inflation Reports, and from then on in the Monetary Policy Reports.⁸ We additionally collect forecasts from the archives of Norges Bank for the period 1994-1996. We use annual forecasts for GDP and CPI since these have the longest coverage and are consistent throughout the sample period.⁹ In 2005, Norges Bank began to base its projections on its own interest rate forecasts. However, that is not a threat to the identification which allows forecasts to respond to expected systematic policy changes.¹⁰

Norges Bank prepared forecasts three or four times a year. Often, the release of these forecasts coincided with a policy meeting. If no forecasts were prepared in conjunction with the policy decision, we assign forecasts iteratively. First, we use forecasts published in the same month of the meeting or in the previous month when possible. Second, for all remaining meetings, we follow [Cloyne and Huertgen \(2016\)](#) in using forecasts by market participants to proxy for the forecasts of the central bank. More precisely, we use data from Consensus Economics, a survey of forecasters from several key private and public-sector institutions, and we resort to the mean of those projections.

For a number of reasons, the forecasts by Consensus Economics are particularly suitable as a proxy. They have a long coverage (starting in 1990), a monthly frequency, are consistent over time, and include annual projections of GDP and the CPI that are directly comparable to the forecasts we obtained from Norges Bank. To ensure that the data from Consensus Economics closely capture Norges Bank's forecasts, we calculate the correlation between both sets of forecasts for periods in which they are both available. Monthly projections of current and next year GDP forecasts have a correlation coefficient of 0.95 and 0.79, and of 0.92 and 0.62 for CPI inflation, respectively. Precisely as with Norges Bank's forecasts, policy meetings are assigned projections from an ear-

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

⁸The historical reports can be found at <https://www.norges-bank.no/en/topics/Monetary-policy/monetary-policy-report/>. The forecasts are typically shown in a table at the end of a report.

⁹We use the GDP forecast for mainland Norway, as opposed to the one for the whole country, for two reasons. First, a larger number of observations of the GDP forecast for mainland Norway are available historically. Second, the Consensus forecast that proxies for missing Norges Bank forecasts, as explained in the text, is only available for mainland Norway.

¹⁰An issue for identification would arise if Norges Bank Staff are aware of non-systematic future policy changes via inside information and would incorporate those expectations into their forecasts. However, we are not aware of instances when this was the case, as Norges Bank Staff only learn about policy decisions when those are officially announced.

lier date in the same month or the previous month.¹¹ Appendix A.2 gives a detailed protocol for the assignment of forecasts to the policy meetings. For a total of 162 policy meetings within our sample period, we associate Norges Bank forecasts to 88 of those and the remaining 74 are filled in using Consensus forecasts.

We exclude the policy meetings that occurred in 1998:M8, 2008:M10, and 2008:M12 from the estimation. The latter two months are associated with meetings that were exceptionally summoned in response to the 2007-09 financial crisis. At this time, macroeconomic projections changed quickly and the forecasts that we have available do not capture those rapid changes. By including these meetings, we would therefore incorrectly assign a larger fraction of the policy rate changes to the shock series, when they were in fact endogenous responses to rapidly changing expectations about the future path of the economy.¹² We also exclude rate changes that occurred in 1998:M8 in response to volatile exchange rate movements and expected exchange rate adjustments, which our controls in (1) may not fully capture, though we find that the results are close to our benchmark when including the 1998:M8 meetings.

2.2 Estimation Results

Based on the described procedure, we estimate (1) using ordinary least squares for the sample 1994:M1-2018:M12. The results are shown in Table 1.

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

Variable	Constant	$i_{m,-1}$	π^m	y^m	$ex_{m,-1}$	$I_{IT} \cdot ex_{m,-1}$
	-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)		
Δ Current Year			0.02 (0.28)	0.27*** (0.00)		
Δ Next Year			0.11** (0.02)	-0.04 (0.58)		
<hr/>						
$N = 162$						
$R^2 = 0.30$						

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

¹¹Regarding the forecast changes, denoted $\Delta\pi_{t+k}^m$ and Δy_{t+k}^m in (1), we use the one relative to the previous forecast for Norges Bank. For the Consensus forecasts, we compare one forecast relative to another two months prior to match the lower frequency of Norges Bank forecasts. For February forecasts, we restrict this gap to one month to increase the number of observations. Various permutations of these assumptions using shorter or longer comparisons did not change our results along relevant dimensions.

¹²For example, for the meetings on 10/15/2008 and 12/17/2008, the latest Norges Bank forecasts were prepared on 6/25/2008 and 10/29/2008, respectively. During such emergency meetings, the central bank likely possesses private information vis-à-vis the market (see, Paul, 2019). Thus, we do not use Consensus forecasts as proxies in these instances.

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.

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, then the shocks are 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 22 and 23 in Appendix A.3.

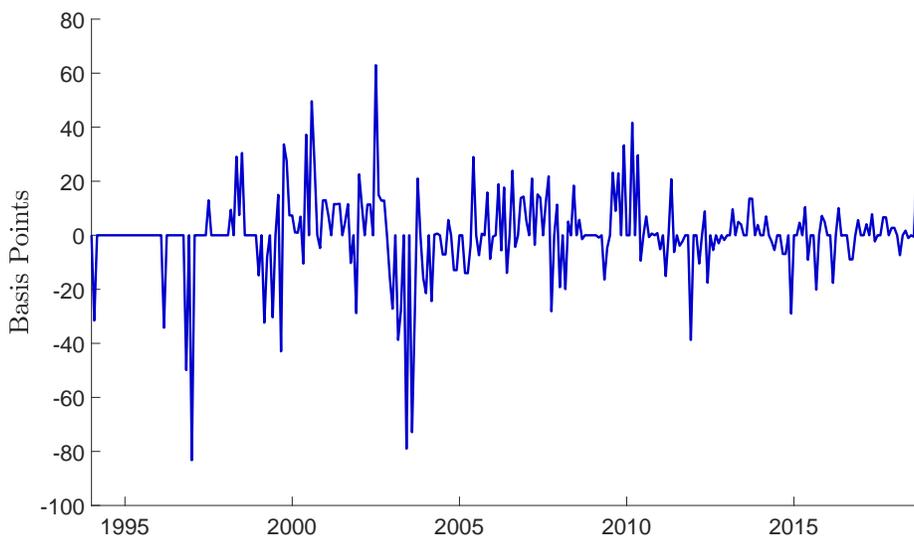


Figure 1: **Monthly Series of Monetary Policy Shocks.**

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 21 in Appendix A.1). [Romer and Romer \(2004\)](#) find similar large movements for the U.S. shock series. Second, the shocks become smaller towards the end of the sample, in line with the reduced volatility of the policy rate in recent years. Third, in Figure 24 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.4 in Appendix A.4).

2.3 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

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

where $h = 0, 1, \dots, 5$ for annual data, $h = 0, 1, \dots, 20$ for quarterly data, and $h = 0, 1, \dots, 60$ for monthly data. The estimated coefficients β^h give the percentage (point) change at horizon h to a 100-basis-point monetary policy shock at the respective frequency.¹³ 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.¹⁴

We consider a wide range of outcome variables at the aggregate level. Table 4 in Appendix A.5 gives precise details on the time series employed.^{15,16} The estimated impulse responses to a contractionary shock of 100 basis points at the monthly, quarterly, and annual frequency are shown in Figures 2, 3, and 4. Across the different frequencies, we obtain textbook responses to a monetary tightening. The policy rate increases, economic activity contracts, as the unemployment rate rises, and industrial production, GDP, and consumption expenditures fall. In addition, consumer prices

¹³Throughout, 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, 2019](#), for details).

¹⁴To 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 Section 2.4. 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.

¹⁵Most 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.

¹⁶The industrial production series shows a structural break over the years considered, which is due to a change in oil extraction, as illustrated in Figure 25 in Appendix A.5. 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.

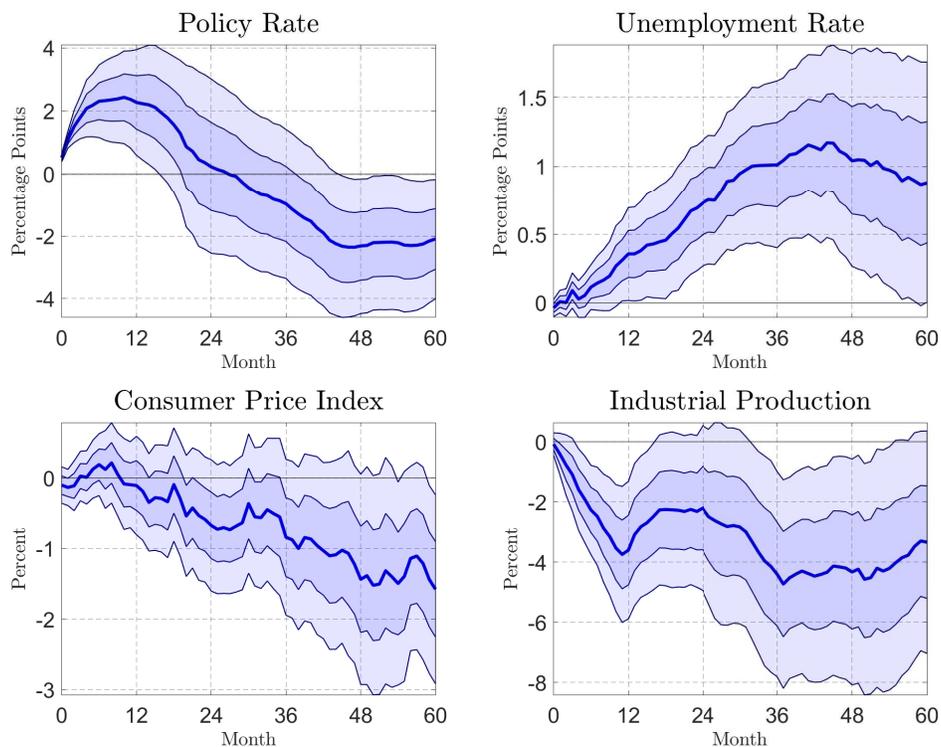


Figure 2: **Impulse Responses at a Monthly Frequency.**

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 4 in Appendix A.5 lists the data sources. Additional impulse responses at a monthly frequency are shown in Figure 26 in Appendix A.6.

and real wages and salaries decline.¹⁷ Figures 26-29 in Appendix A.6 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. Most 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 responses 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.7, 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

¹⁷We 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.

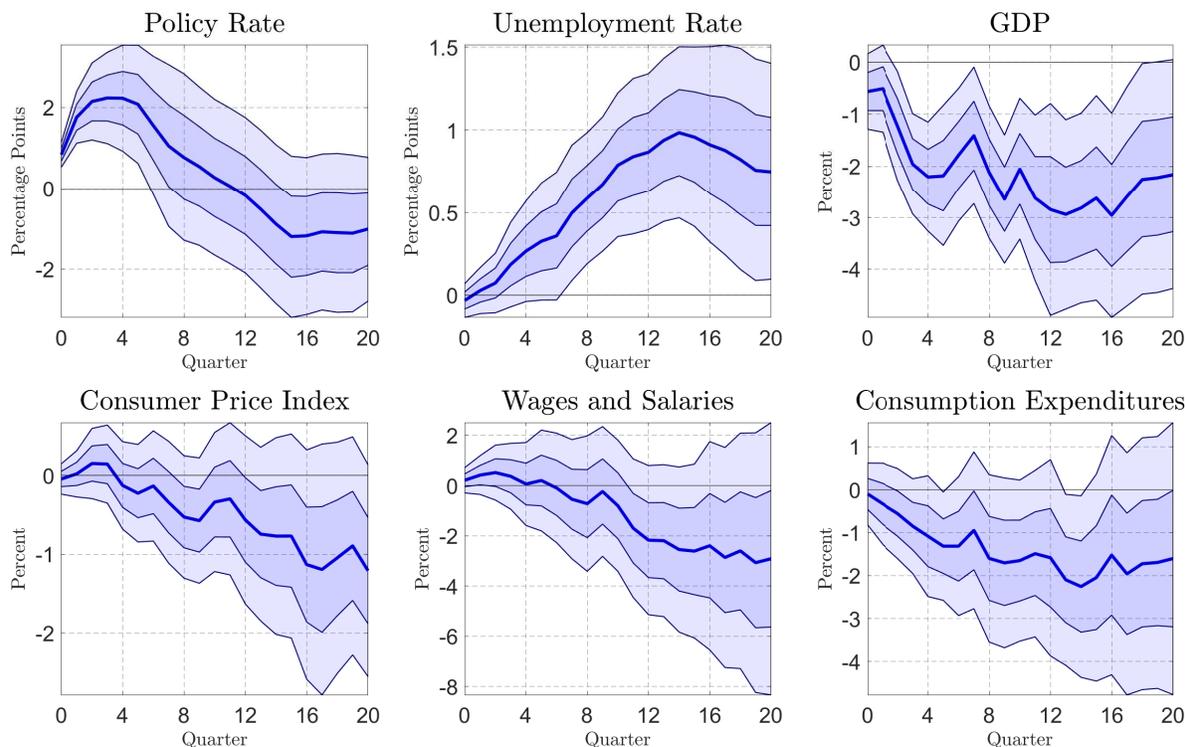


Figure 3: **Impulse Responses at a Quarterly Frequency.**

Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at a quarterly frequency, based on the local projection approach in (2). 95 and 68 percent confidence bands shown, using [Newey and West \(1987\)](#) standard errors. Table 4 in Appendix A.5 lists the data sources. Additional impulse responses at a quarterly frequency are shown in Figure 27 in Appendix A.6.

the shapes and the relative magnitudes of the responses across different frequencies are consistent. Overall, the stability of the impulse responses for various frequencies motivates us to use the annual shocks in combination with the administrative micro data in the following sections. Moreover, at the micro level, we are mainly interested in the heterogeneity of the responses in the cross section of the population and the relative movements of different household balance sheet items. Potential attenuation arising from the time aggregation of the shocks therefore plays a minor role for the analysis that uses household-level data.

2.4 Robustness

Before proceeding to the micro data, we investigate the impulse responses further, compare them to the ones based on U.S. data, and conduct a number of robustness exercises at the macro level. A noticeable feature of our impulse responses, to which we will return when interpreting the micro-responses, is the inverse J-curve reaction of the policy rate. The policy rate initially increases to a contractionary shock, but then falls below its initial level after several months. In Figure 31 in Appendix A.8, we show that this response is inherited from positive (contractionary) shocks. Monetary easings (negative shocks) are not followed by strong subsequent increases. A possible explanation for this result is the following. When the central bank tightens policy beyond the

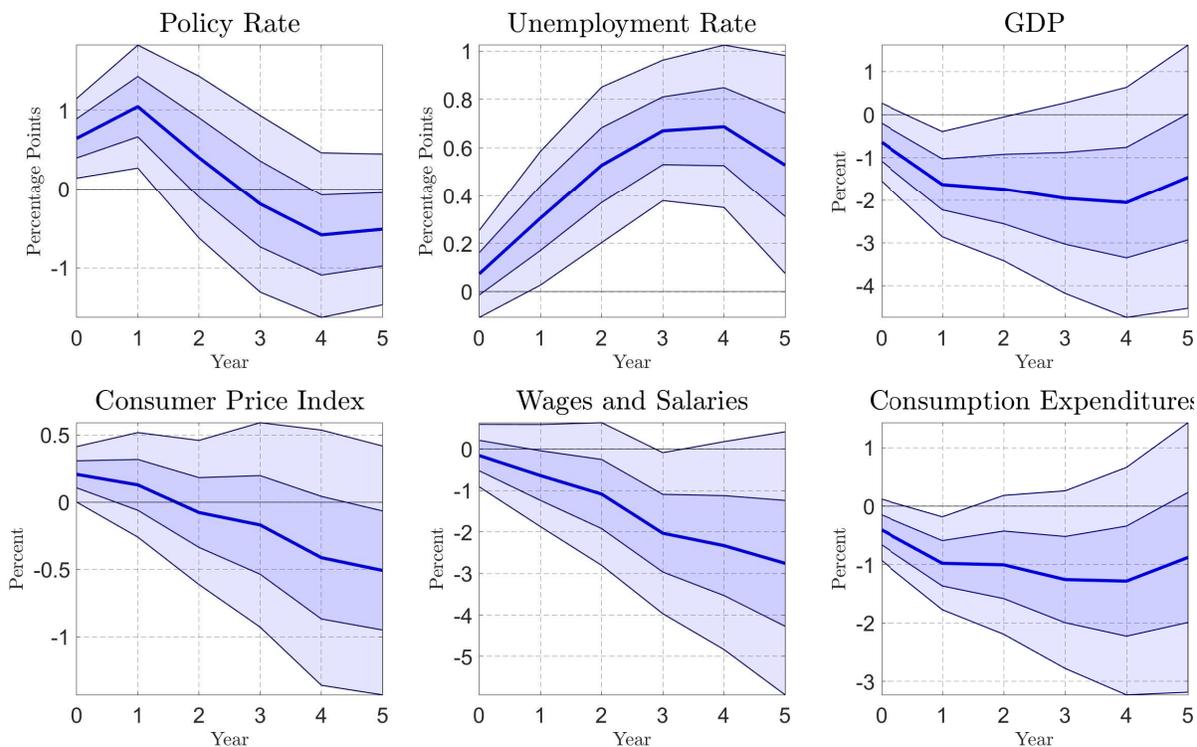


Figure 4: **Impulse Responses at an Annual Frequency.**

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 28 and 29](#) in [Appendix A.6](#).

typical reaction to inflation and output forecasts, policymakers subsequently correct their initial “mistake” and try to boost the economy by easing policy. In contrast, when loosening policy too much, policymakers simply let the boom die out without leaning strongly against it. The estimated responses are close to the ones based on U.S. data, including the inverse J-curve response of the policy rate. [Figure 32](#) in [Appendix A.9](#) shows the responses of the policy rate, industrial production, and the unemployment rate for the United States to monetary policy shocks taken from [Romer and Romer \(2004\)](#) (1970:M1-1996:M12) and [Coibion et al. \(2017\)](#) (1970:M1-2007:M12). [Figure 33](#) in [Appendix A.9](#) shows the policy rate responses to positive and negative shocks for the U.S. Both figures include the equivalent responses for Norway, which tend to be close to the responses based on U.S. data.

We conduct a battery of robustness checks, and the related figures at a monthly frequency are shown in [Appendix A.10](#). First, using only forecasts from Consensus Economics gives similar impulse responses at various frequencies (shown in [Figure 34](#)). Second, the estimated shocks and impulse responses remain nearly unchanged when leaving out the exchange rate terms in (1), but omitting all forecast variables changes the impulse responses significantly. For example, as illustrated in [Figure 35](#), the unemployment rate falls in the wake of a tightening, while wages rise and consumer prices increase. Third, including interaction terms between the output and inflation

forecasts and the pre-inflation targeting regime dummy in (1) gives nearly identical shocks and impulse responses. Fourth, Coibion (2012) notices the importance of the lag length for the results in Romer and Romer (2004). Figure 36 shows that the results are robust to including only one year of lagged values of the monetary policy shock. The same holds for the responses at lower frequencies and when controlling for one year of lagged one-period growth rates of the dependent variable in (2). Fifth, a potential concern related to the aggregation of shocks to annual frequency is that the effects of shocks that take place early in the year are different from those that occur late in the year (e.g., Olivei and Tenreyro, 2007), resulting in biased responses. Figure 37 shows that the results remain similar when considering impulse responses that take place in either the first half or the second half of the year.¹⁸ Finally, as shown in Figure 38, we find nearly identical results when restricting the samples to 1997:M1-2015:M12 or 2001:M1-2018:M12, and the same holds again for the responses at quarterly and annual frequency. Inflation responds slightly less for the sample that starts in 2001, likely because of the official regime switch to inflation targeting in that year.

3 Administrative Data

At the micro level, we base our study on Norwegian administrative data. Norway levies both an income and a wealth tax on its households. The tax authority therefore collects information on all sources of income and balance sheet components down to various asset categories. The data are 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 various sources, the consumption imputation procedure, and further sample restrictions, 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 socio-economic 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 feature several advantages. First, the balance sheets and income statements of the same households are observed across multiple time periods. We are

¹⁸In addition, the policy meetings are fairly evenly distributed throughout the year. Out of the 162 meetings in our sample, 36 take place in the first quarter, 48 in the second quarter, 38 in the third quarter, and 40 in the fourth quarter.

therefore able to construct a panel and follow the responses of households to monetary policy innovations 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, limiting the scope for measurement errors.

3.2 Imputed Consumption Expenditures

We compute a measure of consumption expenditures for each household using the budget constraint:¹⁹

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

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 (wage income, business income), capital income (dividends, interest income, net of interest expenses), transfers (pensions, social security, 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.²⁰ A measure of saving *including* capital gains is directly observed because net wealth is available at the beginning and end of each year at the household level. 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.²¹ For stocks, we use the stockholder registry and capital 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 com-

¹⁹Consumption imputation has by now been widely applied in the literature (see e.g. Leth-Petersen 2010; Eika, Mogstad and Vestad 2020; Fagereng, Holm and Natvik 2018). 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.

²⁰Baker, Kueng, Meyer and Pagel (2018) compare imputed consumption with transaction-level data and show that systematic measurement errors are 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.

²¹Our measure of housing wealth is from Fagereng, Holm and Torstensen (2019). They combine detailed transaction data and information on housing unit characteristics to estimate housing wealth at the household-level using machine learning methods.

pute capital gains. Figure 5 compares the imputed consumption expenditures with consumption in the national accounts. The two consumption series follow each other closely.

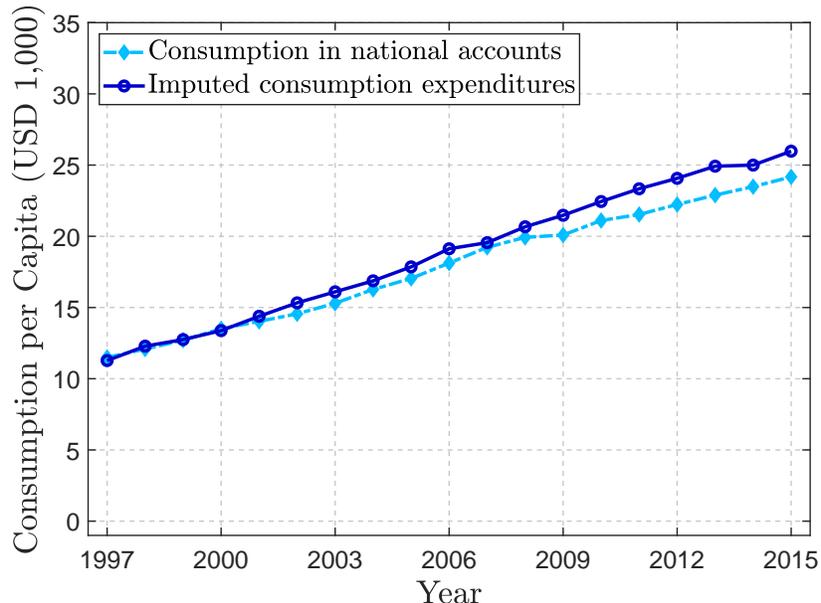


Figure 5: **Imputed Consumption Expenditures and Consumption in the National Accounts.**

Notes: Shown are nominal imputed consumption expenditures based on equation (3) and nominal 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 nominal U.S. dollars.

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 couple and single. Third, we require each individual to have income and consumption above the minimum level in the Norwegian social security scheme.²² And fourth, 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. Table 2 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, and in contrast to the United States where such contracts are typically issued with a fixed rate.²³ In the deposit market, demand deposits account for more than

²²The minimum level in the Norwegian social security was approximately \$11,000 in 2015.

²³In 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 (<https://www.ssb.no/en/bank-og-finansmarked/statistikker/renter>, Table 10648).

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 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 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.

85 percent of all deposit contracts.²⁴ As shown in Figure 39 in Appendix B.2, overall deposit rates closely track the policy rate. In the United States, deposit rates of 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).

A potential concern is that the data do not cover information on pension wealth. However, the absence of data on pension wealth is unlikely to affect our results for the following reason. The main share of pension wealth of Norwegian households consists of defined benefits associated with the public pension system. As households work, they accumulate pension points that are translated into pension income when they retire. The pension income that they receive is independent of the performance of financial markets and thus monetary policy. Financial risk is instead born by the Norwegian government, which holds substantial wealth in the Norwegian Government Pension Fund Global ("the oil fund").

²⁴Between 1996 and 2015, the share of deposits with adjustable rates was between 85 and 92 percent (<https://www.ssb.no/en/bank-og-finansmarked/statistikker/banker>, Table 11003).

3.5 Income, Wealth, and Liquid Assets

Following the recent literature on heterogeneous-agent modeling, we initially focus on the effects of monetary policy for 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.²⁵ 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 6 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 7 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.²⁶ 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.

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 (since they are net debtors) in the group containing the median household to more than 55 percent in the top 0.1 percent group.

²⁵This 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.

²⁶The 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.

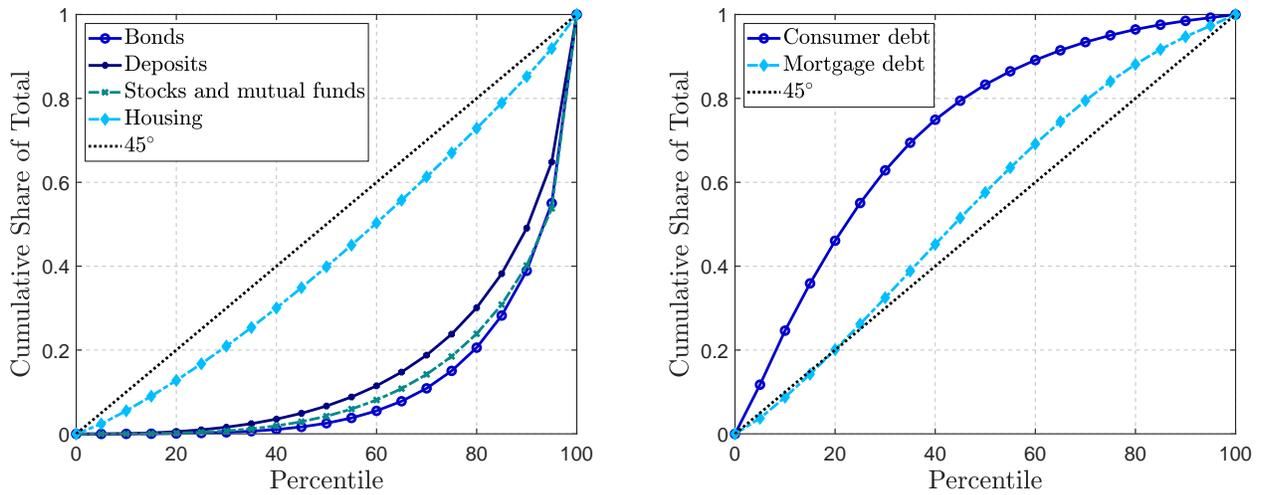


Figure 6: **Cumulative Shares of Assets and Liabilities.**

Notes: Cumulative shares by asset and debt type against percentiles of liquid asset distribution.

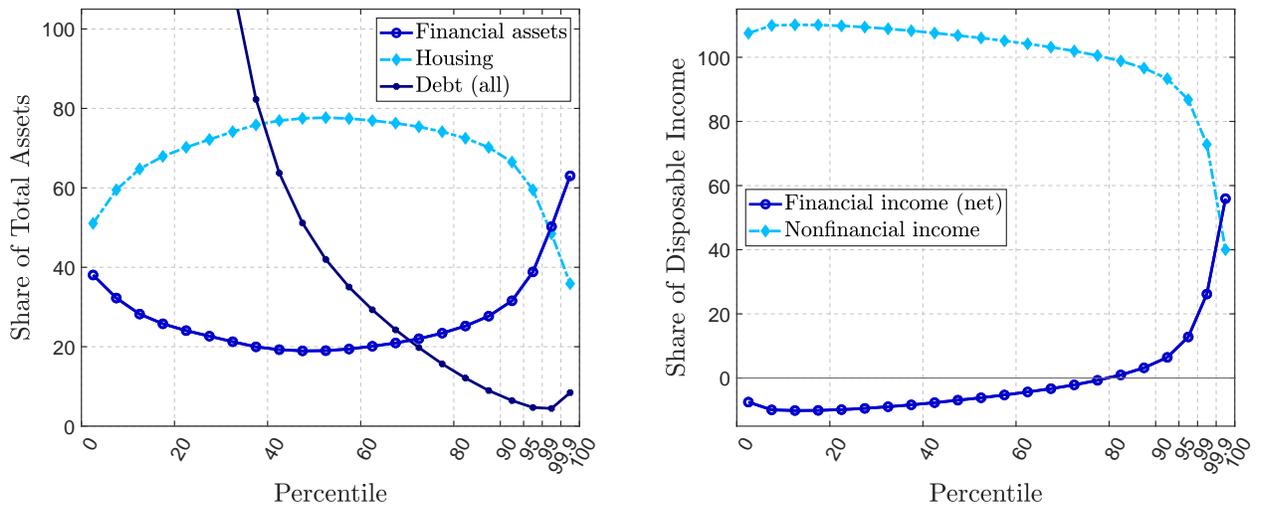


Figure 7: **Portfolio and Cash-Flow Shares.**

Notes: Shares of total assets (left panel) and disposable income (right panel) against percentiles of liquid asset distribution. Group averages shown.

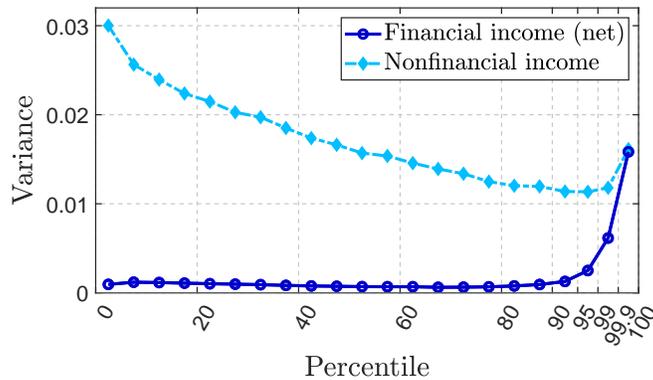


Figure 8: **Income Volatility.**

Notes: Variance of income growth components weighted by income shares. Group averages shown. See text and footnote 27 for details.

To the extent that household consumption is responsive to income changes and the response of financial income to monetary policy is distinct from that of nonfinancial income, we should observe a differential consumption response at different segments of the liquid asset distribution even in the absence of heterogeneity in MPCs. The same is true if the response of a given income type varies across the distribution. Figure 8 reinforces these points. It portrays the average variance of financial and nonfinancial income growth weighted by the corresponding income share for the same groups as above.²⁷ At the bottom of the distribution, households are exposed to a significant amount of income risk through their nonfinancial income, which predominantly consists of labor earnings. The income volatility incurred through financial income is negligible. As initially non-financial income volatility declines with liquid assets and financial income volatility remains low, households around the 90th percentile of the distribution face comparably low fluctuations from both individual sources of income. At the upper end of the distribution, the volatility of financial income is drastically elevated. For the top 0.1 percent, it is of the same magnitude as the volatility of nonfinancial income. The following section shows that heterogeneity in the income response to monetary policy lies at the heart of differences in the consumption response among households located at both ends of the liquid asset distribution.

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

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

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-

²⁷More precisely, since disposable income $inc_{i,t}$ is the sum of financial income $inc_{i,t}^f$ and nonfinancial income $inc_{i,t}^{nf}$, disposable income growth is given by $g_{i,t} = (inc_{i,t-1}^f / inc_{i,t-1})g_{i,t}^f + (inc_{i,t-1}^{nf} / inc_{i,t-1})g_{i,t}^{nf}$ where $g_{i,t}^x$ with $x \in \{f, nf\}$ is the growth rate of the respective income type. The figure shows the time variance of both parts of the sum, averaged across all households in a given group. Only households observed for at least five consecutive years enter the computation. Note that the variance of $g_{i,t}$ is the sum of the two curves plotted and a covariance term.

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.²⁸

Notably, the dependent variable is defined as the change in y_i from $t - 1$ to $t + h$, normalized by the average value \bar{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 (2.3). Conditional on the other covariates, β^h gives the average h -period change of y_i across all households, in units of one-period lagged per-capita \bar{y}_{t-1} , in response to a monetary policy shock.²⁹ Standard errors which are robust to heteroskedasticity, autocorrelation, and cross-sectional dependence are calculated as laid out in [Driscoll and Kraay \(1998\)](#).

The estimated impulse responses are shown in Figure 9 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 4. Initially they increase, and subsequently fall. Their response suggests that they are driven by interest income and expenditures, channels that we investigate in more detail below. Further, household debt, wealth, and risky assets fall, whereas the response of safe assets is slightly positive.³⁰ Overall, the responses go in the expected directions, and they are statistically different from zero at the 95 percent confidence interval, or at least at the 68 percent confidence interval.

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 40 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 15 below).

²⁸Specifically, we include $\frac{y_{i,t-1}-y_{i,t-2}}{y_{i,t-2}}$ and $\frac{y_{i,t-2}-y_{i,t-3}}{y_{i,t-3}}$ 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.

²⁹In 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.

³⁰Household 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 [Fagereng, Holm, Moll and Natvik, 2019](#)). Risky 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.

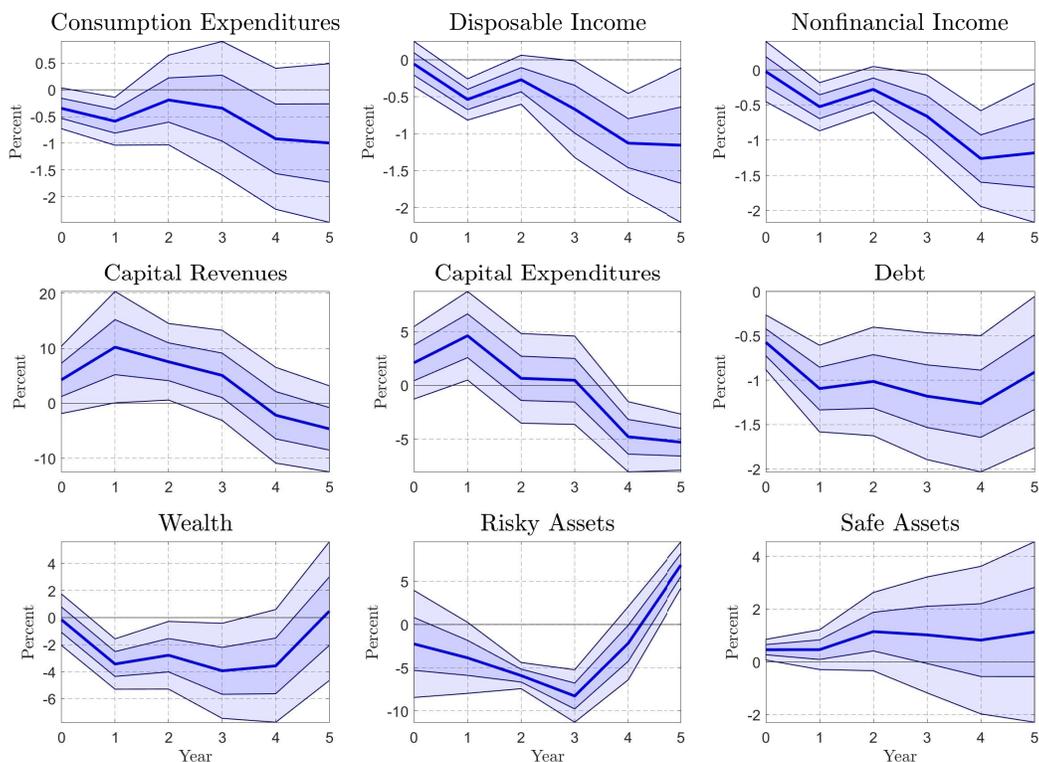


Figure 9: **Micro-Macro Impulse Responses.**

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.

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 $g = 1, 2, \dots, 10$ 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 occurring in period t . For each group g , the setup of the local projections is

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

where the notation is as in Section 4.1 and $h = 0, 1, \dots, 5$. 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 $inc_{i,t-1}$.³¹ 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.

³¹Changes in the consumption-to-income ratio are not normalized by lagged income.

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 exhibits 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. We refer to consumption out of capital gains as “wealth effects.”

4.2.1 Consumption, Income, and Saving

Figure 10 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 11. 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.³² For each variable, the figures portray the complete impulse responses of all 10 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 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 41-42 show the full impulse responses and the associated confidence intervals for selected groups ($g = 1, 5, 10$) and Figures 43-44 the relative responses across groups, which allow to evaluate the significance of group differences. Overall, the dynamics and the group differences highlighted below are statistically different from zero at the 68 or 95 percent confidence level.

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 average decline in years 4 and 5 following the shock is about 1.4 percent. On impact, consumption and the consumption-to-income ratio decline while saving increases somewhat. In the following years, the saving response is small compared to the income response, so that consumption closely follows income. Consumption of the households located

³²Throughout, we distinguish between *saving* (flow) and *savings* (stock).

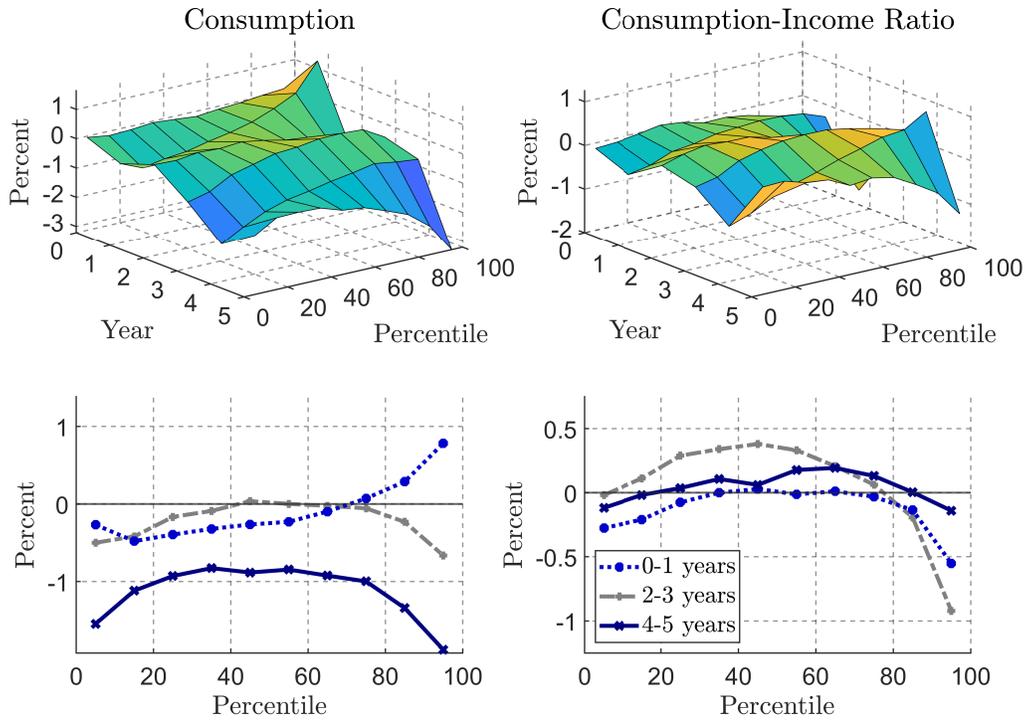


Figure 10: **Impulse Responses of Consumption and Consumption-Income Ratio.**

Notes: Changes relative to lagged net income (left) and changes of ratio (right) at segments of the liquid asset distribution. See Figures 41-44 in Appendix C.2 for group responses and group comparisons with confidence bands.

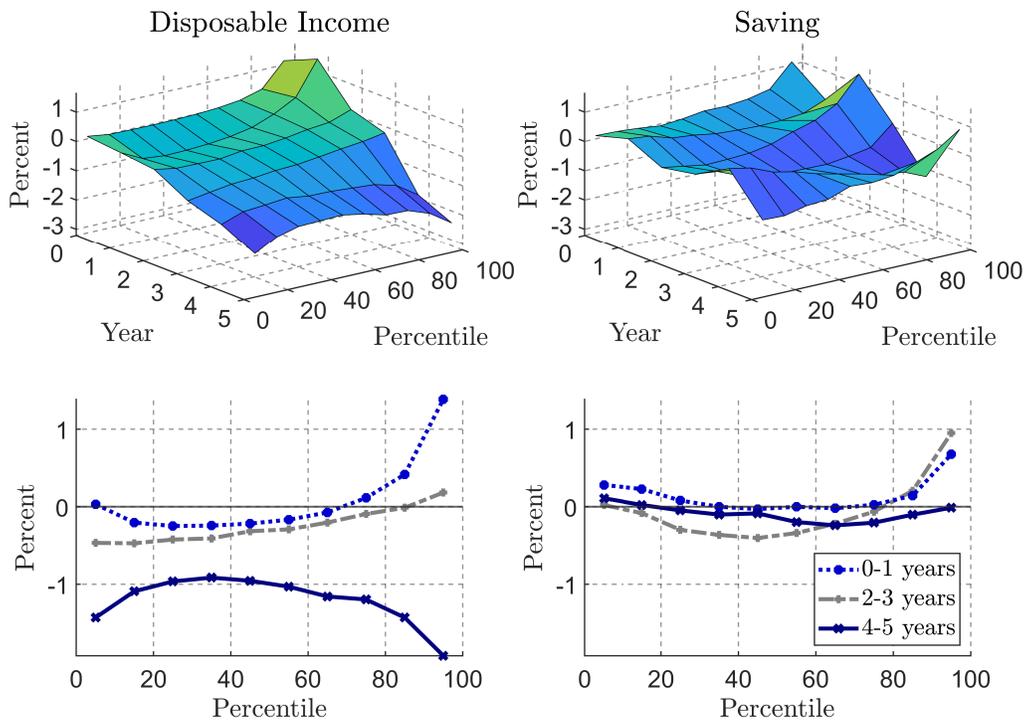


Figure 11: **Impulse Responses of Net Income and Saving.**

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures 41-44 in Appendix C.2 for group responses and group comparisons with confidence bands.

in the middle of the distribution is less responsive on average. In conjunction with their income decline in years 2 and 3 after the shock, they allow their consumption-to-income ratio to rise and hence saving to fall. In years 4 and 5, consumption is lowered but, relative to initial household income, the size of the decline is only a bit more than half of the response of the bottom 10 percent of households.

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 roughly the same amount. Income falls sharply following the initial increase. Because the decline in income is not matched by a sizable adjustment of the saving behavior, consumption also falls. While households at the top of the distribution show an active saving response on impact, the income change feeds into a 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 marginally larger but of similar magnitude as that at the bottom of the distribution. As a result, the consumption response is inverse U-shaped along the liquid asset distribution.

Several of the results are worth stressing. As can be seen from Figure 10, the consumption response is delayed, with the peak occurring four to five years after the shock in all groups. At the bottom, the peak response is significantly larger than around the median and the response of the consumption-to-income ratio is comparably small and stable across horizons. Both of these observations are consistent with the HANK literature, although the mechanisms at work may differ. The model by [Kaplan, Moll and Violante \(2018\)](#), for example, 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 nearly unaffected. The lack of a significant saving response is consistent with hand-to-mouth behavior at these horizons. Since initially saving increases and consumption falls somewhat when the policy rate rises, the estimates do not allow us to reject that intertemporal substitution is of importance 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.³³

Our findings for the top of the distribution do not support the predictions of the HANK model in [Kaplan, Moll and Violante \(2018\)](#). In their model, the households with the largest amounts of liquidity holdings documented have small MPCs and their responses to monetary policy are primarily driven by the substitution effect of the policy rate change. If this was also the case in

³³The 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.

our data, the estimates would show an increase in saving and a fall in consumption following the initial monetary tightening. The large increase in consumption for high-liquidity households can also not be explained by a standard income effect resulting from the change in lifetime income associated with the rise in the interest rate. In Appendix D, we illustrate this fact using a canonical consumption-saving model with a borrowing constraint and a liquid financial asset. In the model, the consumption of agents with very large liquidity holdings and a low elasticity of intertemporal substitution rises following a contractionary monetary policy shock. However, even if the size of the income effect relative to the substitution effect is exaggerated by selecting an unrealistically low elasticity of intertemporal substitution, the peak consumption response is orders of magnitude too small compared to the empirical estimates (see Figure 64). Thus, our results suggest that the households at the top of the distribution have sizable MPCs and that their behavior is not predominantly driven by standard intertemporal substitution motives. The estimates are consistent instead with cash-flow effects playing an important role.³⁴ We elaborate on this point in the following section, where we investigate the effect of monetary policy on household income in more detail.

4.2.2 Decomposing the Income Response

Below, we decompose the income responses portrayed in Figure 11 into various subcomponents. A first set of results is contained in Figure 12. 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 gradually falls after a small initial increase until an average decline of about 2.1 percent of disposable income is reached in years 4 and 5. Net financial income is lowered by 0.3 percent of lagged disposable income 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 is smaller, amounting to 1.3 percent, but net financial income initially increases by 2.3 percent of household income before a decline of 1.6 percent is reached. Hence, the majority of the hump-shaped response of disposable income after four to five years is driven by nonfinancial income for low-liquid-asset households and by financial income for high-liquid-asset households.

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 7 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

³⁴Consistent with this view, Fagereng, Holm and Natvik (2018) find MPCs out of lottery prizes that are substantially larger than zero even for households with high levels of liquidity. Figure 67 in Appendix E.4 shows MPC estimates, which are derived using our series of identified monetary policy shocks. Both sets of estimates line up closely.

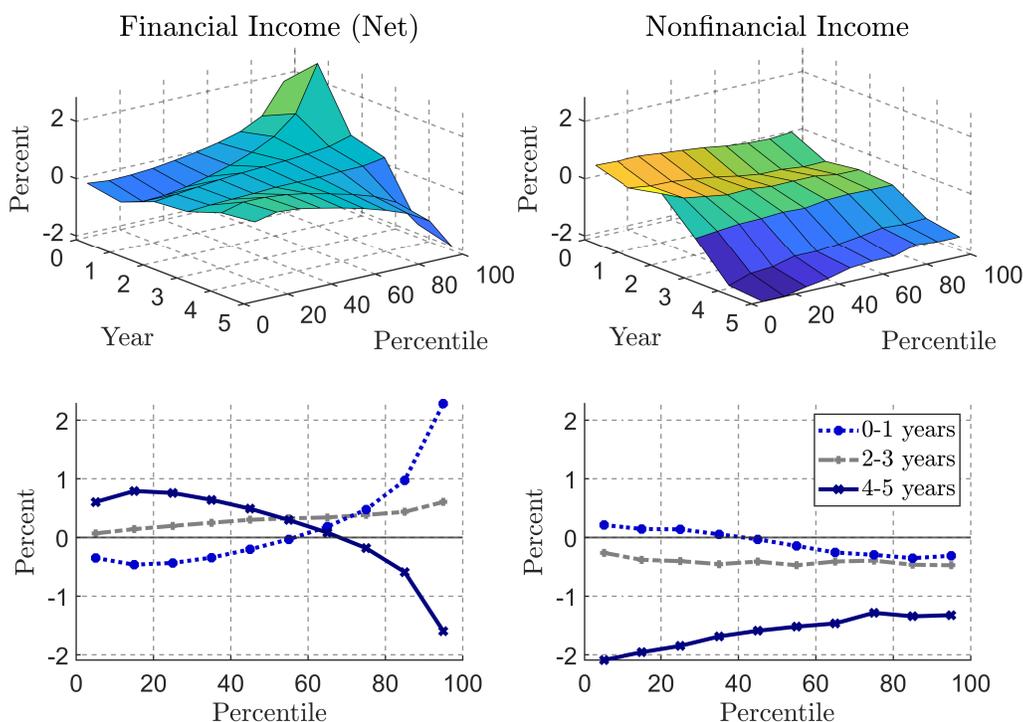


Figure 12: **Impulse Responses of Net Financial Income and Nonfinancial Income.**

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures 41-44 in Appendix C.2 for group responses and group comparisons with confidence bands.

bottom. Correspondingly, the nonfinancial income response is stronger 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 13. 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 14 shows the impulse responses of two capital-revenue components, interest income and dividend income. Note that dividend income is 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 \(1956\)](#).

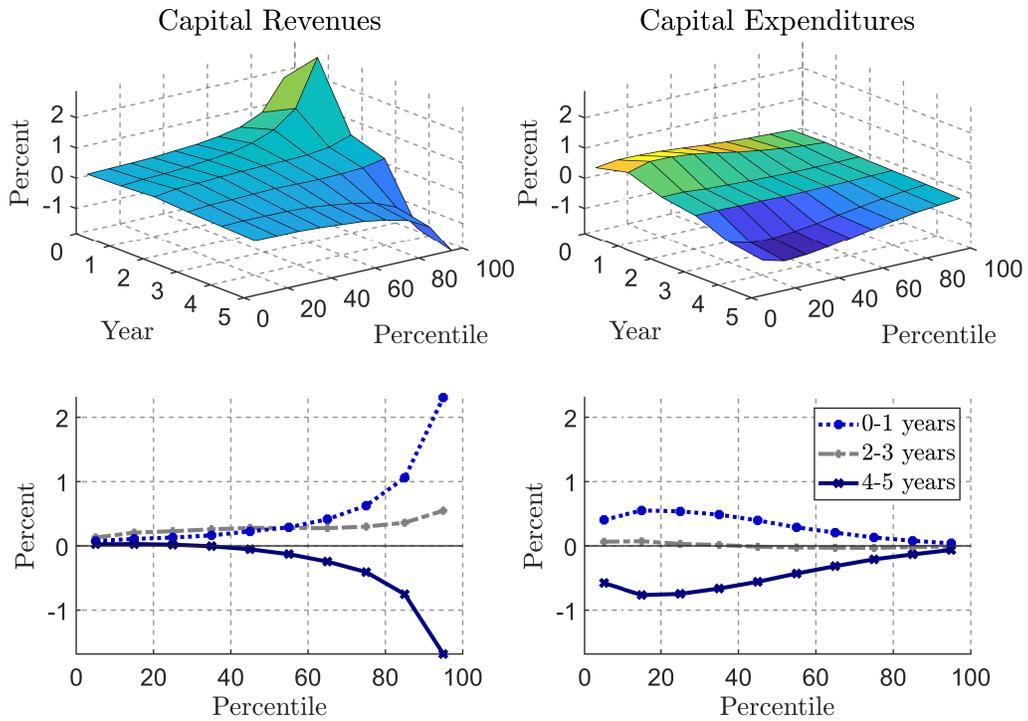


Figure 13: **Impulse Responses of Capital Revenues and Expenditures.**

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures 41-44 in Appendix C.2 for group responses and group comparisons with confidence bands.

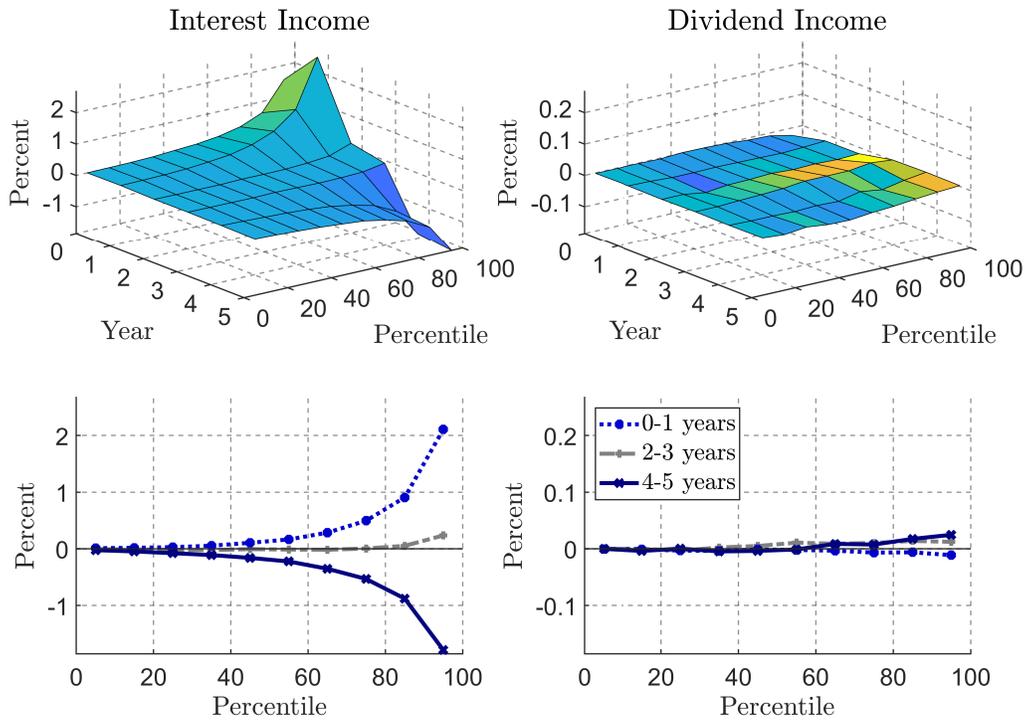


Figure 14: **Impulse Responses of Interest and Dividend Income.**

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures 41-44 in Appendix C.2 for group responses and group comparisons with confidence bands.

Figures 12-14 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 income or substitution 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 \(2020\)](#), 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. Investment adjustment costs imply that investment reacts sluggishly and stickiness in household expectations implies that consumption responds at the same time as labor income. [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 Wealth Effects

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 assets in Figure 15. 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 ([Di Maggio, Kermani and Majlesi, 2020](#); [Fagereng et al., 2019](#); [Chodorow-Reich, Nenov and Simsek, 2019](#)). The estimates reflect that housing is more equally distributed than risky 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 in units of lagged disposable income are therefore more 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 45 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, Holm and Natvik \(2018\)](#) as is shown in Figure 46. Even households at the top of the liquid asset distribution may be affected by borrowing constraints if their income and consumption is large

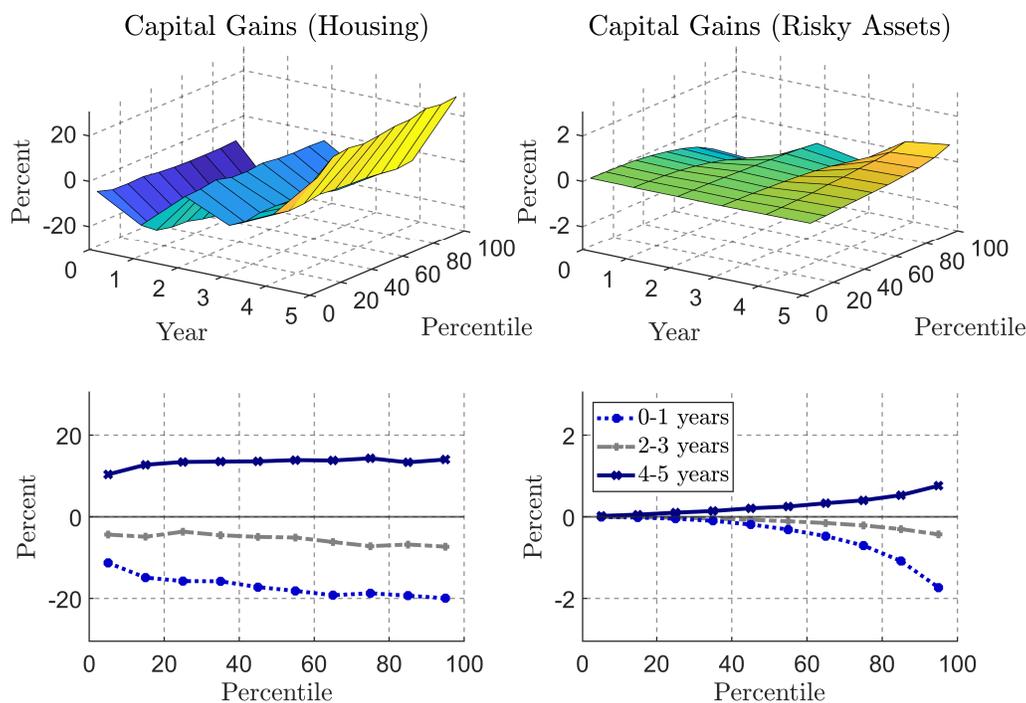


Figure 15: **Impulse Responses of Capital Gains in Housing and Risky Assets.**

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. See Figures 41-44 in Appendix C.2 for group responses and group comparisons with confidence bands.

compared to their liquidity holdings. To exclude this possibility, Figure 47 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 in stock holdings does not drive the estimates, we exclude stockholders from the estimation. The results, plotted in Figure 48, are nearly unchanged again. Finally, Figures 49 and 50 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 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 fall into the top 10 percent is substantial, higher than USD 120,000 in some years. Second, the large majority of households in the top group remain among the top 10 percent every year.³⁵ The high persistence in liquid asset holdings makes it unlikely that households with large planned expenditures self-select into the top of the distribution.

³⁵On 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.

4.3 Which Households Drive the Aggregate Effects of Monetary Policy?

The magnitude of the responses to monetary policy shocks in units of lagged household income may differ along the liquid asset distribution due to heterogeneity in both the change in the variable of interest or the level of income prior to the shock. To isolate the heterogeneity of the former type, the one associated with active adjustments made to consumption, for example, we re-estimate equation (5) replacing $inc_{i,t-1}$ with average disposable income across all households. In this way, the units are made comparable in absolute terms across liquid asset groups. More broadly, we are interested in whether household groups exist that are of disproportionate importance for the aggregate responses. The model estimated is given by

$$\frac{y_{i,t+h} - y_{i,t-1}}{\overline{inc}_{t-1}} = \delta_i^h + \beta_g^h \cdot \epsilon_i^{MP} + \sum_{k=1}^K \gamma_{g,k}^h X_{i,t-k} + u_{i,t}^h \quad \forall i \in g, \quad (6)$$

where $\overline{inc}_t = (\sum_i inc_{i,t})/N_t$ and N_t is the number of households contained in the sample in year t . The remaining variables are defined as before.

The results are shown in Figure 16. Analogous to Section 4.2, a detailed set of results with confidence intervals is contained in Appendix C.4. Several of the patterns described in the previous section remain visible since disposable income is relatively equally distributed across liquid asset groups.³⁶ As shown in Section 4.2, households at the bottom exhibit a larger consumption response relative to their own income than those around the mid-point. Figure 16 illustrates that low-liquidity households do not make a substantially larger contribution to the aggregate response. The average response in years 4 and 5 following the shock is remarkably stable across all groups below the 80th percentile. In contrast, the average consumption response in the top group is more than twice as large as that of the group containing the median household ($g = 5$). Thus, households with large liquid asset holdings have a disproportionately strong influence on aggregate consumption.

As the upper 10 percent of households hold more than half of all interest rate-sensitive financial assets, their capital income response is significantly larger than that of any other group. In combination with a nonfinancial income response of similar size four to five years after the shock, disposable income in the top group falls almost three times as much as in the groups surrounding the median household. Because debt is concentrated at the bottom of the distribution, one might expect households with few liquid assets to experience large net financial income changes as well. The figure shows that this is indeed the case. However, the changes in capital expenditures at the bottom are of smaller magnitude than the changes in capital revenues at the top. Figure 6 shows that this can be explained by the fact that the concentration of debt at the bottom is smaller than the concentration of deposits and bonds at the top.

³⁶Average disposable income across all years ranges between \$34,841 and \$53,916 in the group with the lowest and highest liquid asset holdings, respectively. The average across all groups is \$43,488. Hence, according to a simple back-of-the-envelope calculation, one percentage point in Figure 16 is roughly equivalent to \$435.

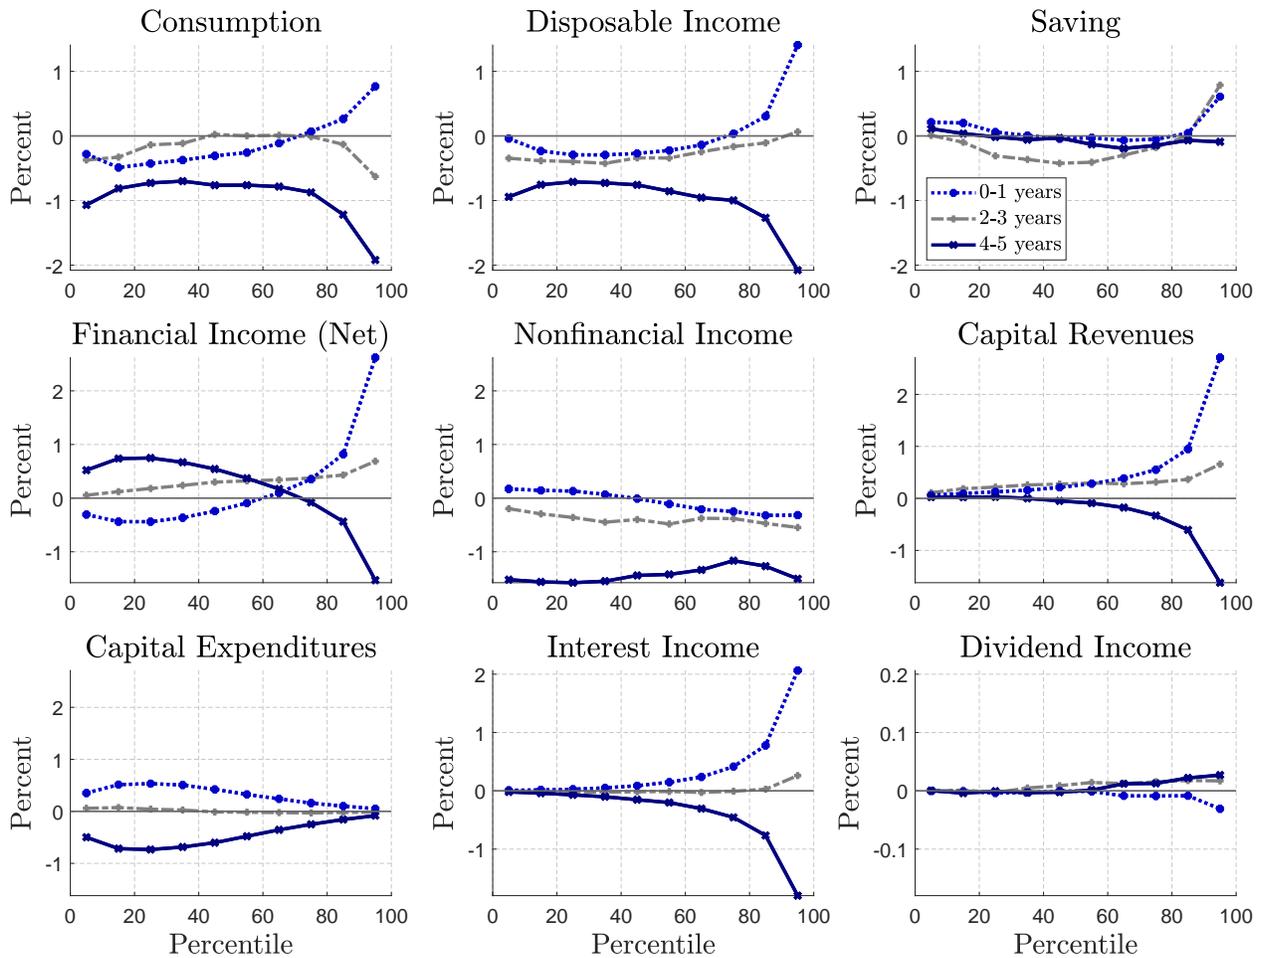


Figure 16: **Impulse Responses in Units of Average Net Income.**

Notes: Changes relative to average lagged net income along the liquid asset distribution. See Figures 51-54 in Appendix C.4 for group responses and group comparisons with confidence bands.

Since the households located above the 90th percentile of the liquid asset distribution contribute most to the aggregate effect of monetary policy, we disaggregate the estimates in this group further. Figure 55 in Appendix C.4 shows the results from estimating the local projections in (6) for the three subgroups formed by the households located between the 90th and 95th, the 95th and 99th, and above the 99th percentile. For nearly all variables, there is a clear ordering of the response magnitudes among the subgroups that is a continuation of the patterns observed along the entire distribution. The top 1 percent of households shows particularly strong consumption, disposable income, and saving responses, for example. However, the differences within the top decile are too small for the top 1 percent to exert a large impact on the total response.

4.4 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 re-estimating the local projections in (5). Changes in the out-

come variables are scaled again in units of each household’s disposable income. 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. Households at the bottom of the distribution are net debtors while households at the top are net creditors. 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 17 and 18. The full impulse responses together with statistical significance tests can be found in Appendix C.5.

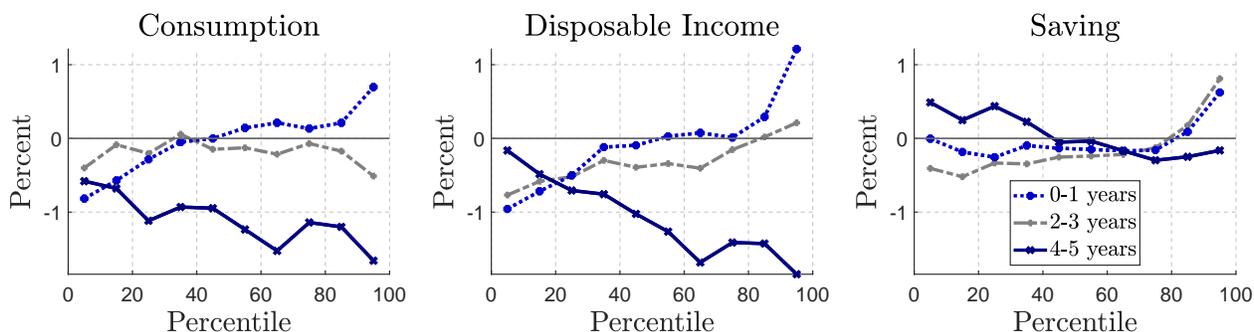


Figure 17: **Impulse Responses by Net Interest Rate Exposure.**

Notes: Changes relative to lagged net income at segments of the distribution of net interest rate exposure. See Figures 60-63 in Appendix C.5 for group responses and group comparisons with confidence bands.

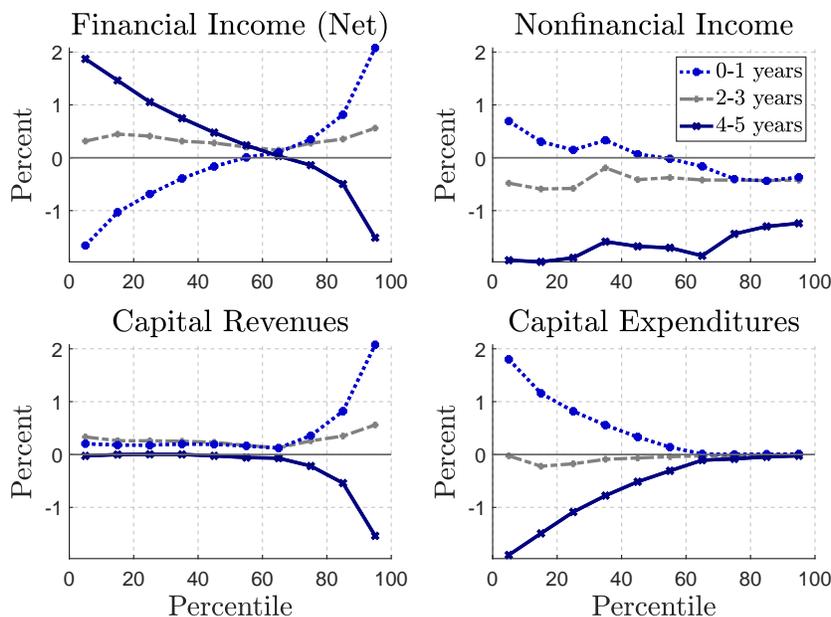


Figure 18: **Impulse Responses of Income Components by Net Interest Rate Exposure.**

Notes: Changes relative to lagged net income at segments of the distribution of net interest rate exposure. See Figures 60-63 in Appendix C.5 for group responses and group comparisons with confidence bands.

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 of opposite sign. The

mutated consumption response at the aggregate level, visible in Figure 9, is a result of opposing effects at the micro level. Households with high positive net interest rate exposure tend to hold large amounts of interest rate-sensitive assets but little or no debt. Correspondingly, their capital revenues increase sharply in response to a positive interest rate shock while their capital expenditures remain unaffected. The strong reaction of net financial income overcompensates a small decline in nonfinancial income, implying that disposable income rises at the top of the distribution.

At the other end of the spectrum, 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 a mirror image of the one by households at the top of the distribution. The large decline in financial income resulting from a strong 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-to-one with disposable income at the bottom.

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 negative across the entire distribution, as above. At the bottom, the decline in capital expenditures is nearly offset by the decline 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 in which debt-servicing costs are low. At the top, the financial and nonfinancial income responses are both negative, implying that disposable income strongly declines. While the saving response also turns negative, consumption falls significantly more than at the bottom of the distribution.

Our results are in line with the findings of several papers 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 U.S. households with adjustable-rate mortgages in the aftermath of the Great Recession. They show that the 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. \(2017\)](#) equally find evidence for interest rate changes to feed into household consumption through adjustable-rate mortgages. While these papers focus on borrowers, [La Cava, Hughson and Kaplan \(2016\)](#) estimate the elasticity of household consumption to interest-sensitive cash flows separately for borrowers and lenders using a survey of Australian households. Concentrating on contemporaneous effects, [La Cava, Hughson and Kaplan \(2016\)](#) find that borrowers display a larger elasticity than lenders. Our results indicate that, relative to their respective lagged income, lenders react more strongly to monetary policy shocks than borrowers, albeit with a time lag. At the 4-5 year impulse response horizon, financial and nonfinancial income decline for net lenders, jointly exerting a contractionary effect on consumption.

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

$$\frac{c_{i,t+h} - c_{i,t-1}}{\bar{c}_{t-1}} = \delta_i^h + \beta^h \epsilon_t^{MP} + \sum_{m=1}^3 \mu_m^h X_{i,t-m} + \sum_{k=1}^K \sum_{m=0}^h \gamma_m^{h,k} \tilde{y}_{i,t+m}^k, \quad (7)$$

for $h = 0, 1, \dots, 5$. 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 (7) is the additional set of income controls $\sum_{k=1}^K \sum_{m=0}^h \gamma_m^{h,k} \tilde{y}_{i,t+m}^k$, where $\tilde{y}_{i,t+m}^k = (y_{i,t+m}^k - y_{i,t-1}^k) / \bar{c}_{t-1}$ is a particular subcomponent k of disposable income.³⁷ 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 (7) 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 (7) 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 (7) 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 (7) are associated with indirect effects of monetary policy.

³⁷All 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.

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 in empirical work, arising from the omission of the income controls in (4). To identify the direct effect using (7), two conditions have to be satisfied. First, households' consumption changes cannot feed back into their income within the same period, otherwise (7) 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, households' income needs to move due to structural variation that does not also affect households' consumption directly, unless it enters as a control in (7). Idiosyncratic labor 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 labor 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 (7), since the data only cover realized income at the household level and not household expectations of future income paths. Second, 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 9 and, e.g., [Christiano, Eichenbaum and Evans, 2005](#)).³⁸

The results are shown in Figure 19.³⁹ 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 4) 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 after three years, which can be associated with the fall in the policy rate around this time (see Figure 4). At the five-year horizon, indirect effects dominate, accounting for more than

³⁸That 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.

³⁹For 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.

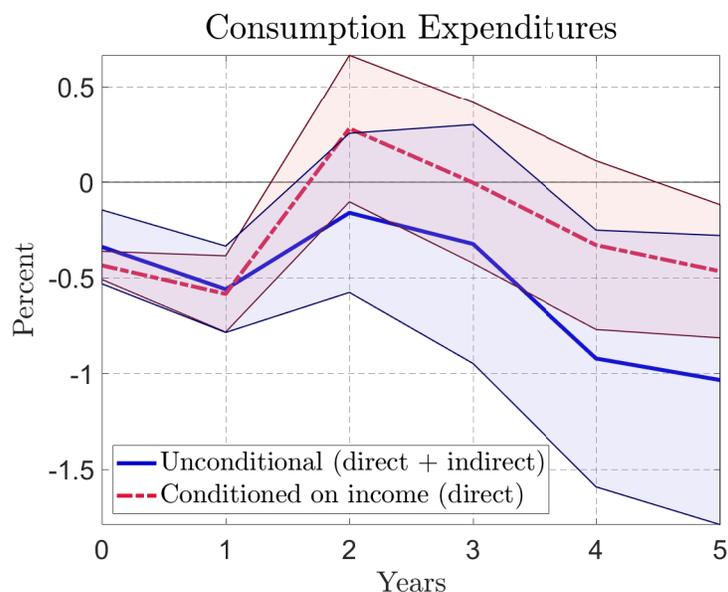


Figure 19: **Direct and Indirect Effects of Monetary Policy.**

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 (7). 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.

50 percent of the consumption response.⁴⁰

Thus, our findings show that the dominance of the indirect effect of monetary policy, introduced 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 66 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 decile, indirect effects explain nearly all of the drop in consumption.⁴¹

A potential concern for identification using regression (7) is that household income may be driven by structural variation that also affects consumption directly and that is not controlled for in (7).

⁴⁰In comparison, [Kaplan, Moll and Violante \(2018\)](#) estimate that indirect effects account for around 80 percent of the total consumption response over the first year.

⁴¹Based on these regressions, Figure 67 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 (7) 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 MPC estimates from [Fagereng, Holm and Natvik \(2018\)](#) for different levels of liquidity, as shown in Figure 67.

If such variation is large, regression (7) 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 \(2018\)](#). 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 (7). In the second stage, the predicted values are used as controls as in (7). Figure 20 shows the results. Strikingly, the estimated impulse responses in Figures 19 and 20 are nearly the same.

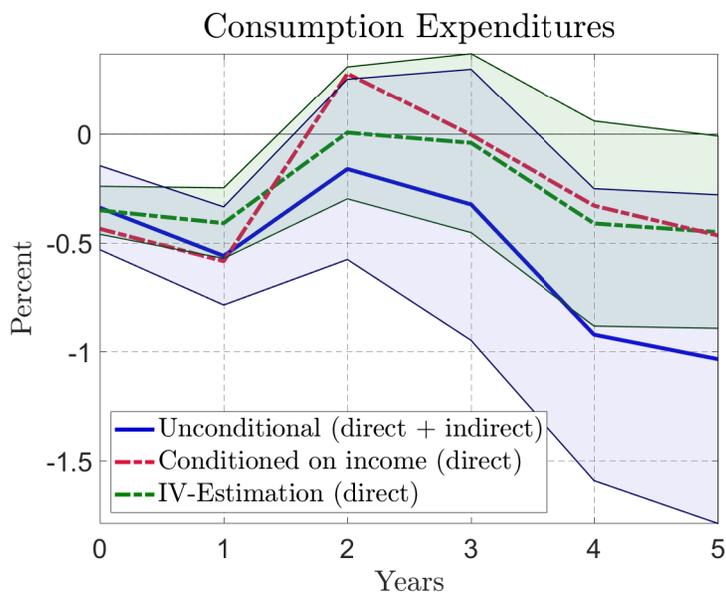


Figure 20: **Direct and Indirect Effects of Monetary Policy – IV-Estimation.**

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 (7) (red dashed line). 68 percent confidence bands shown, using [Driscoll and Kraay \(1998\)](#) standard errors.

In Figure 68 in Appendix E.5, we provide an alternative decomposition of the consumption response into direct and indirect effects using MPC estimates from [Fagereng, Holm and Natvik \(2018\)](#) and the disposable income response shown in Section 4.1. We obtain similar results: initially the direct effect dominates, but after several years the indirect effect takes over. From year two onwards, the indirect plays an even larger role compared with Figure 19, explaining nearly all of the consumption response.

6 Conclusion

Until recently, empirical monetary economics largely focused on the response of aggregate time series to monetary policy innovations. A common theme throughout our paper is that such aggregate reactions mask vast heterogeneity in the responses across households underneath the surface. In particular, we obtain the following results that have several implications for the new generation of HANK models.

First, by separating households according to their liquid asset holdings, we find that households with few liquid assets show a stronger response of consumption and income compared with households that hold a typical amount of liquid assets. Following a monetary tightening, these households tend to lower their consumption-to-income ratio, whereas the one of the median household rises. These results suggest that financial frictions that hinder consumption smoothing play a key role for households with few liquid assets – a mechanism that is at the heart of the vast majority of recent HANK models.

Second, monetary policy changes the net financial income of borrowers and savers, and we find that these changes largely feed through into household consumption. When the policy rate rises, net creditors see their interest income increase, whereas net borrowers experience a rise in their interest expenses. We find that both of these income changes substantially impact household consumption, but with the opposite sign. Hence, these results show the importance of cash-flow channels of monetary policy operating through household net interest income. To account for these channels, several target moments should be matched within HANK models: (i) the distribution of assets and debt holdings for households, (ii) the sensitivity of interest income and expenses to policy rate changes, and (iii) marginal propensities to consume out of changes in interest income and expenses depending on households' asset and debt positions.

And third, we provide the first complete empirical estimate of the relative importance of direct and indirect effects, following the distinction made by [Kaplan, Moll and Violante \(2018\)](#). We find that the significance of these two effects depends on the impulse response horizon. In the first two years after a shock, direct effects of monetary policy dominate. In the following years, indirect effects outweigh direct effects, while the response of nonfinancial income builds up. Our results therefore confirm the key finding in [Kaplan, Moll and Violante \(2018\)](#) that indirect effects of monetary policy play a larger role – with the important difference that we show that it takes time for these effects to take over. A fruitful avenue for future work would be to replicate our finding based on an extension of the model by [Kaplan, Moll and Violante \(2018\)](#).

References

- Ampudia, Miguel, Dimitris Georgarakos, Jiri Slacalek, Oreste Tristani, Philip Vermeulen and Giovanni L. Violante. 2018. Monetary policy and household inequality. Working Paper 2170 European Central Bank.
- 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.
- 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." *The 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. 2019. "The New Keynesian cross." *Journal of Monetary Economics*, forthcoming.
- Bilbiie, Florin O., Diego R. Känzig and Paolo Surico. 2020. Capital, Income Inequality, and Consumption: The Missing Link. Unpublished manuscript.
- 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.
- Chodorow-Reich, Gabriel, Plamen T Nenov and Alp Simsek. 2019. Stock market wealth and the real economy: A local labor market approach. Working Paper 25959 National Bureau of Economic Research.
- 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, Clodo Ferreira and Paolo Surico. 2019. "Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism." *The 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.
- Debortoli, Davide and Jordi Gali. 2018. Monetary Policy with Heterogeneous Agents: Insights from TANK models. Unpublished manuscript.
- 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." *forthcoming, Journal of Finance* .
- Drechsler, Itamar, Alexi Savov and Philipp Schnabl. 2017. "The Deposits Channel of Monetary Policy." *The 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." *The 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 Vestad. 2020. "What can we learn about household consumption from information on income and wealth." *forthcoming, Journal of Public Economics* .
- 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, 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. 2018. MPC Heterogeneity and Household Balance Sheets. Unpublished manuscript.
- Fagereng, Andreas, Martin B. Holm and Kjersti N. Torstensen. 2019. Housing wealth in Norway, 1993-2015. Unpublished manuscript.

- Floden, Martin, Matilda Kilstrom, Josef Sigurdsson and Roine Vestman. 2017. Household Debt and Monetary Policy: Revealing the Cash-Flow Channel. CEPR Discussion Papers 12270.
- Gertler, Mark and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7(1):44–76.
- 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." *The Quarterly Journal of Economics* 132(3):1427–1467.
- 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.
- Kuester, Keith, Nils Gornemann and Makoto Nakajima. 2016. Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy. 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.
- La Cava, Gianni, Helen Hughson and Greg Kaplan. 2016. The Household Cash Flow Channel of Monetary Policy. RBA Research Discussion Papers rdp2016-12 Reserve Bank of Australia.
- 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. 2019. Transmission of Monetary Policy with Heterogeneity in Household Portfolios. Unpublished manuscript.
- 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.
- 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. 2019. "The Time-Varying Effect of Monetary Policy on Asset Prices." *Review of Economics and Statistics*, forthcoming.
- Ravn, Morten and Vincent Sterk. 2018. Macroeconomic Fluctuations with HANK and SAM: an Analytical Approach. Unpublished manuscript.
- 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. Working Paper 26630 National Bureau of Economic Research.
- 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.

A Monetary Policy Identification

A.1 Historical Policy Rate

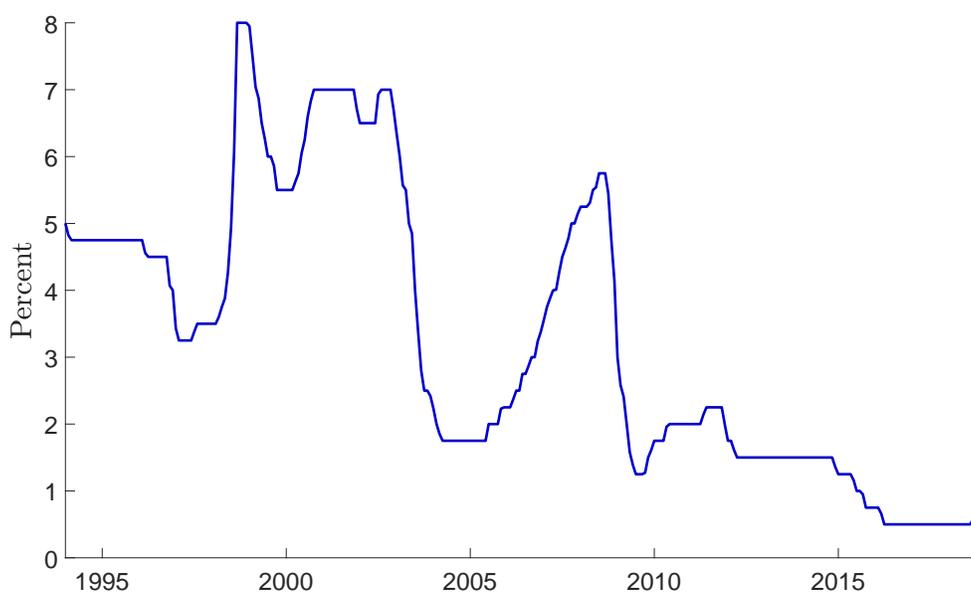


Figure 21: Historical Policy Rate (Sight Deposit Rate).

A.2 Protocol for Forecast Assignment

To assign forecasts to the policy meetings (162 in total), we use the following rule:

1. If available, use Norges Bank forecasts that are either
 - directly prepared for a policy meeting (51),
 - the same month before the meeting (5), or
 - in the month before the meeting (32).
2. For any remaining meetings, we use the Consensus forecasts that are
 - conducted in the same month before the meeting (4), or
 - in the month before the meeting (70).

Hence, we use the Norges Bank forecasts for 88 meetings, and the Consensus forecasts for the remaining 74 meetings.

A.3 Monetary Policy Shock Series

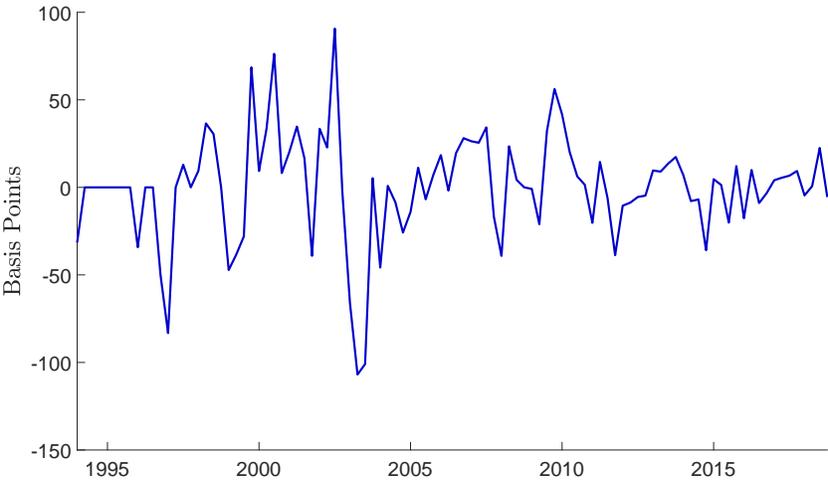


Figure 22: Quarterly Series of Monetary Policy Shocks.

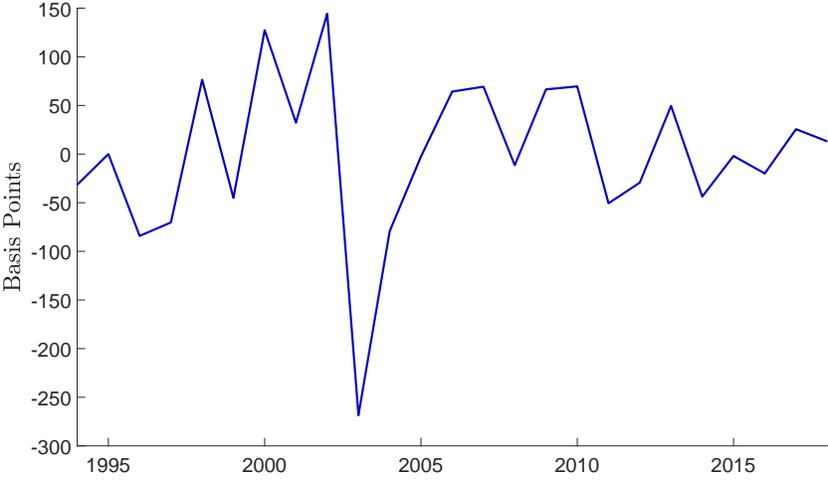


Figure 23: Annual Series of Monetary Policy Shocks.

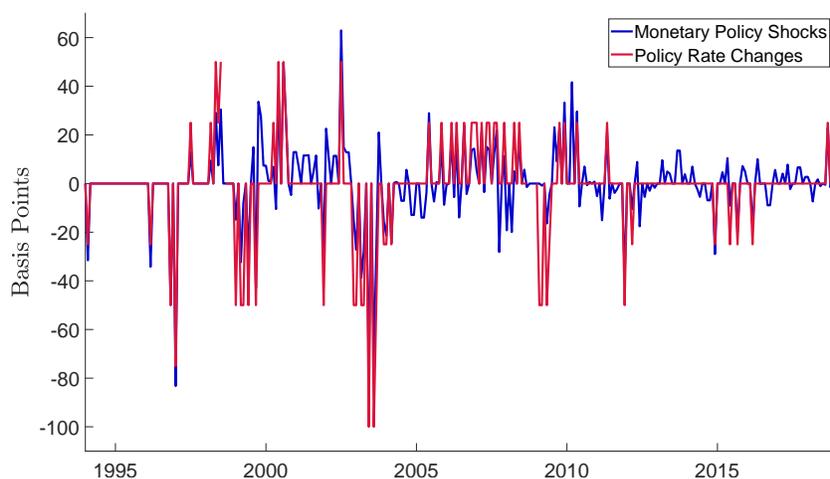


Figure 24: Monetary Policy Shocks vs. Policy Rate Changes.

A.4 Predictability of Monetary Policy Shocks

Table 3: Predictability of Monthly Monetary Policy Shocks.

Variables	3 Lags		6 Lags		9 Lags	
	F-statistic	P-value	F-statistic	P-value	F-statistic	P-value
Unemployment Rate	0.90	0.44	0.83	0.55	0.60	0.80
CPI-AEL Inflation	1.15	0.33	1.42	0.21	1.36	0.21
Industrial Production	0.25	0.86	0.25	0.96	0.27	0.98
All of the above	0.84	0.58	0.86	0.63	0.73	0.84

Notes: The dependent variable is the monthly series of monetary policy shocks. The regressors are three, six, or nine lagged values of the change in the unemployment rate, monthly CPI-AEL inflation, the monthly growth rate of industrial production, or a joint regression with all three variables. The table reports F-statistics and the associated p-values given the null hypothesis that all coefficients are zero.

A.5 Macro Data

Table 4: Macro Time Series.

Variable Name	Annual	Quarterly	Monthly	Start	Source
Policy Rate: Sight Deposit Rate	x	x	x	1993	Norges Bank
Unemployment Rate: Registered, SA	x	x	x	1995	FRED
Consumer Price Index: AEL, All-Item Index	x	x	x	1995	Statistics Norway
Industrial Production: Smoothed Index, SA	x	x	x	1993	Statistics Norway
Manufacturing Production: Index, SA			x	1993	FRED
Producer Prices Manufacturing: Index			x	1993	FRED
Exchange Rate: Norway Kroner/U.S. Dollar Exchange Rate			x	1993	FRED
GDP: Real GDP, COP, SA	x	x		1993	FRED
GDP Deflator: Implicit Price Deflator Norway, Index, SA	x	x		1993	FRED
Real Gross Fixed Capital Formation: COP, SA	x	x		1993	Statistics Norway
House Prices: Price index for existing dwellings, CPI-ADJ.	x	x		1993	Statistics Norway
Wages and Salaries: Mainland Norway, CPI-ADJ., SA	x	x		1995	Statistics Norway
Consumption Expenditures: COP, SA	x	x		1993	Statistics Norway
Hours Worked: Employees & self-employed, Mainland, SA	x	x		1995	Statistics Norway
Household Income: After-tax income, COP	x			1993	Statistics Norway
Top 10% Income Share: After-tax income distribution	x			1993	Statistics Norway
Gini coefficient: After-tax income distribution	x			1993	Statistics Norway
P90/P10: After-tax income distribution	x			1993	Statistics Norway
Extraction and Related Services: SA			x	1993	Statistics Norway
Average Loan Rate	x			1993	Statistics Norway
Average Deposit Rate	x			1993	Statistics Norway
U.S. Data					
Policy Rate: Effective Federal Funds Rate			x	1970	FRED
Unemployment Rate: Civilian, SA			x	1970	FRED
Industrial Production: Index, SA			x	1970	FRED

Notes: “SA” stands for “Seasonally-adjusted”, “CPI-ADJ.” indicates that the series was converted into real units using the consumer price index, “COP” stands for “Constant Prices”. Monthly, quarterly, and annual indicate the frequency for which the impulse responses are shown. “Start” indicates the year of the earliest observations that are used. While the analysis starts in 1993, time series may be available beforehand.

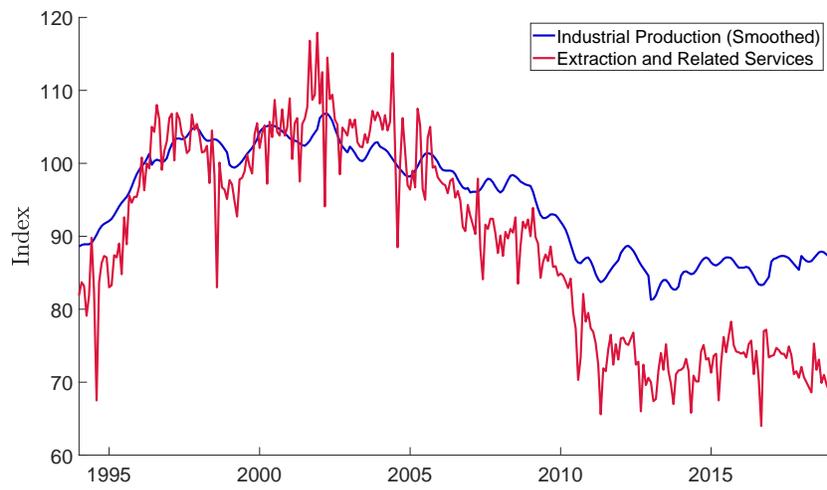


Figure 25: Structural Break in Industrial Production.

Notes: Table 4 in Appendix A.5 lists the data sources.

A.6 Additional Impulse Responses

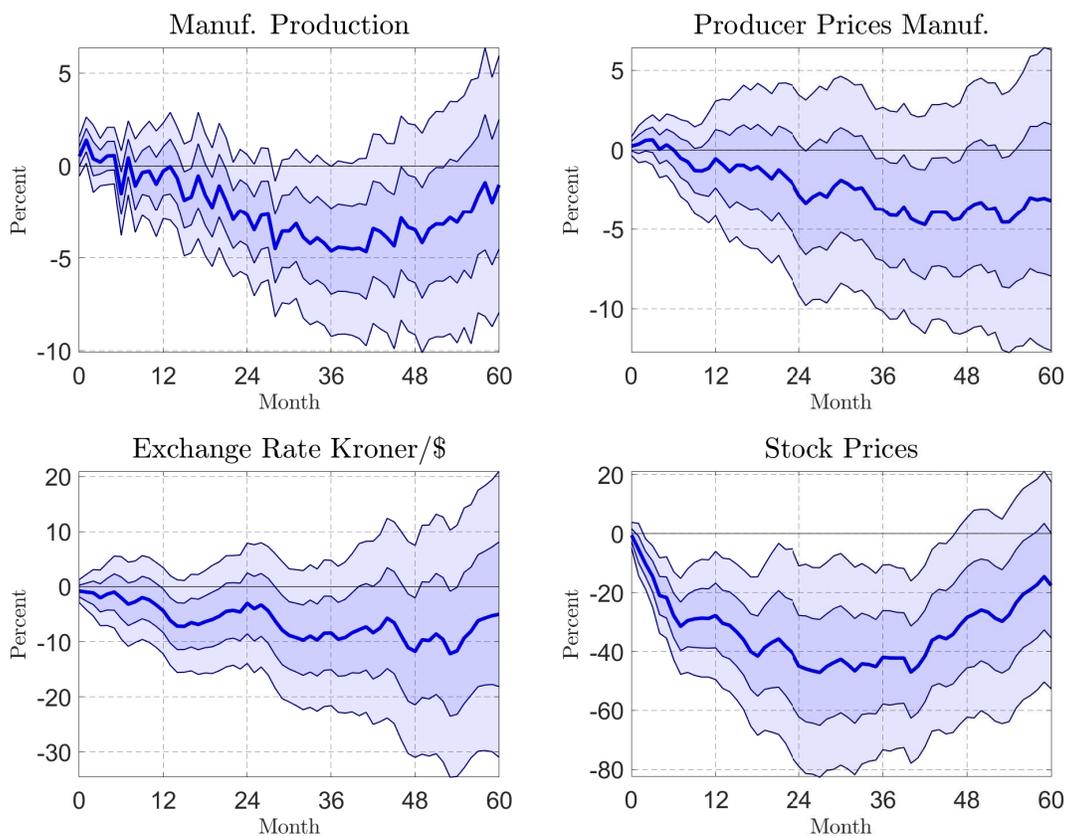


Figure 26: **Impulse Responses at a Monthly Frequency.**

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 4 in Appendix A.5 lists the data sources.

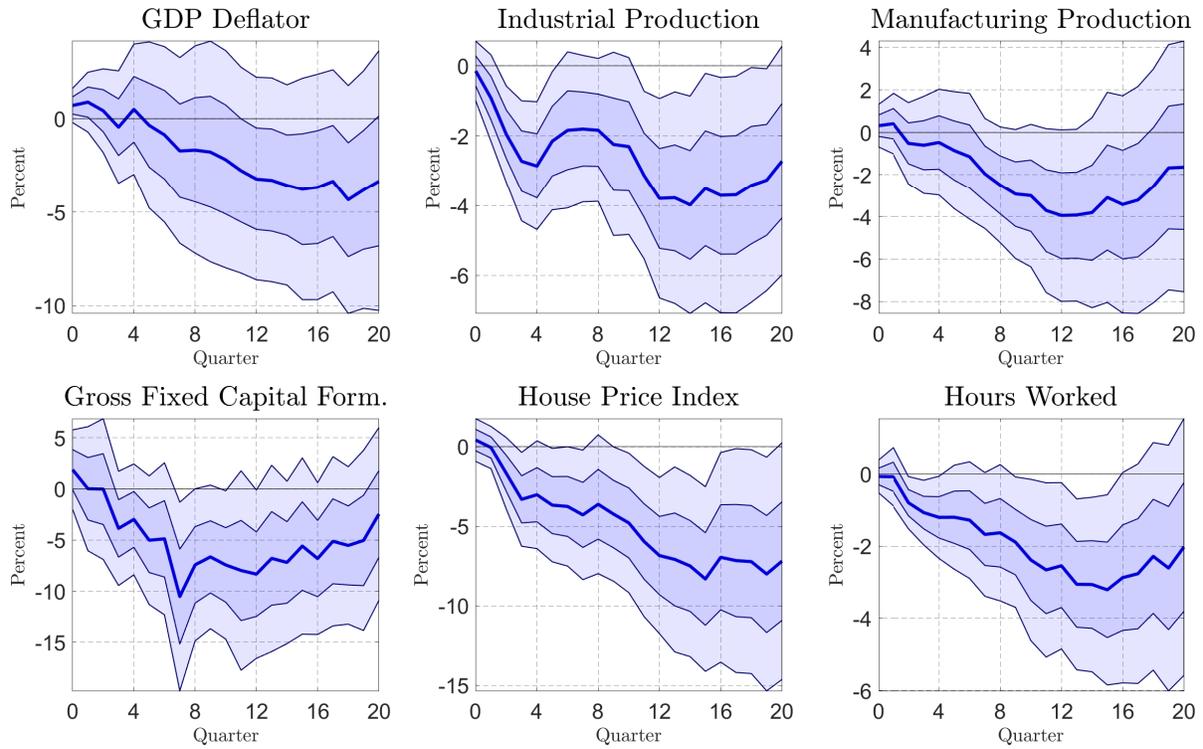


Figure 27: **Impulse Responses at a Quarterly Frequency.**

Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at a quarterly frequency, based on local projection approach in (2). 95 and 68 percent confidence bands shown, using [Newey and West \(1987\)](#) standard errors. Table 4 in Appendix A.5 lists the data sources.

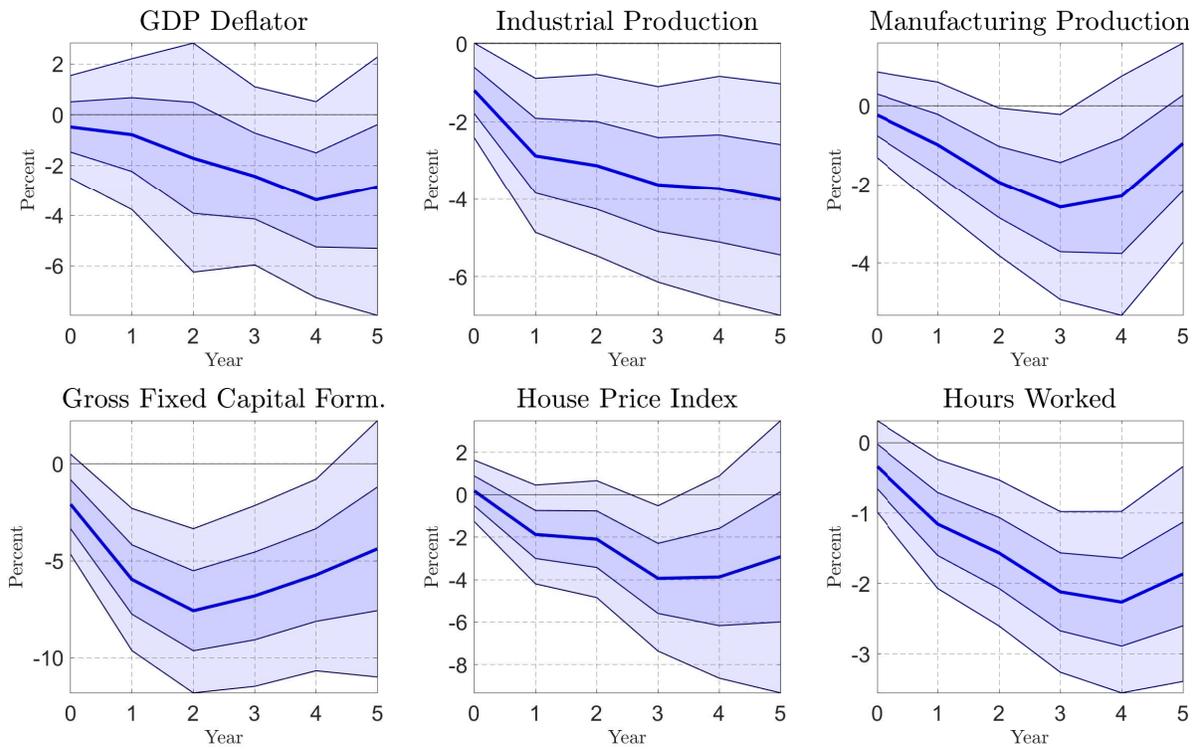


Figure 28: Impulse Responses at an Annual Frequency.

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. Table 4 in Appendix A.5 lists the data sources.

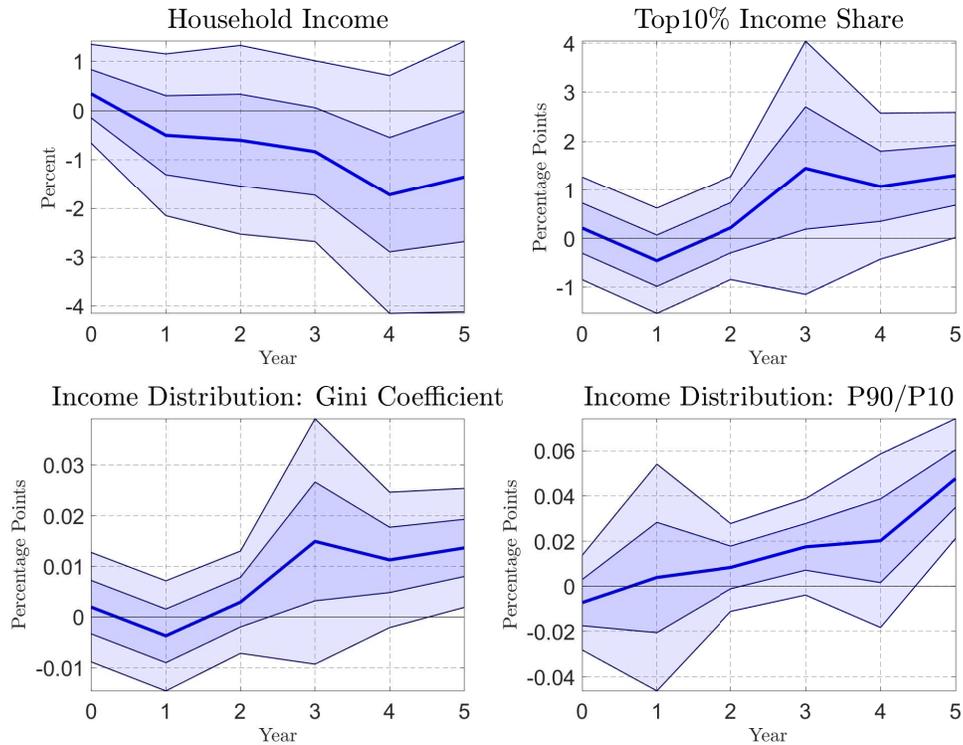


Figure 29: **Impulse Responses of Income Inequality Measures.**

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. Table 4 in Appendix A.5 lists the data sources.

A.7 Time Aggregation

Figures 2-4 show that the impulse responses at an annual frequency are generally of smaller magnitude compared with the ones at monthly and quarterly frequencies. To test whether the size of the responses is still similar in relative terms though, we proceed as follows. Using the estimated monthly response of a particular variable, we aggregate this response to the annual frequency. To this end, we assume that one observes the monthly response equally distributed throughout the year. Hence, this procedure takes the estimated monthly response as the true response and computes the corresponding reaction at the annual frequency. The results for several variables that are observed at the monthly frequency are shown in blue in Figure 30 as “annualized monthly” responses. In comparison, the red lines indicate the “estimated annual” responses, which are of smaller magnitude.

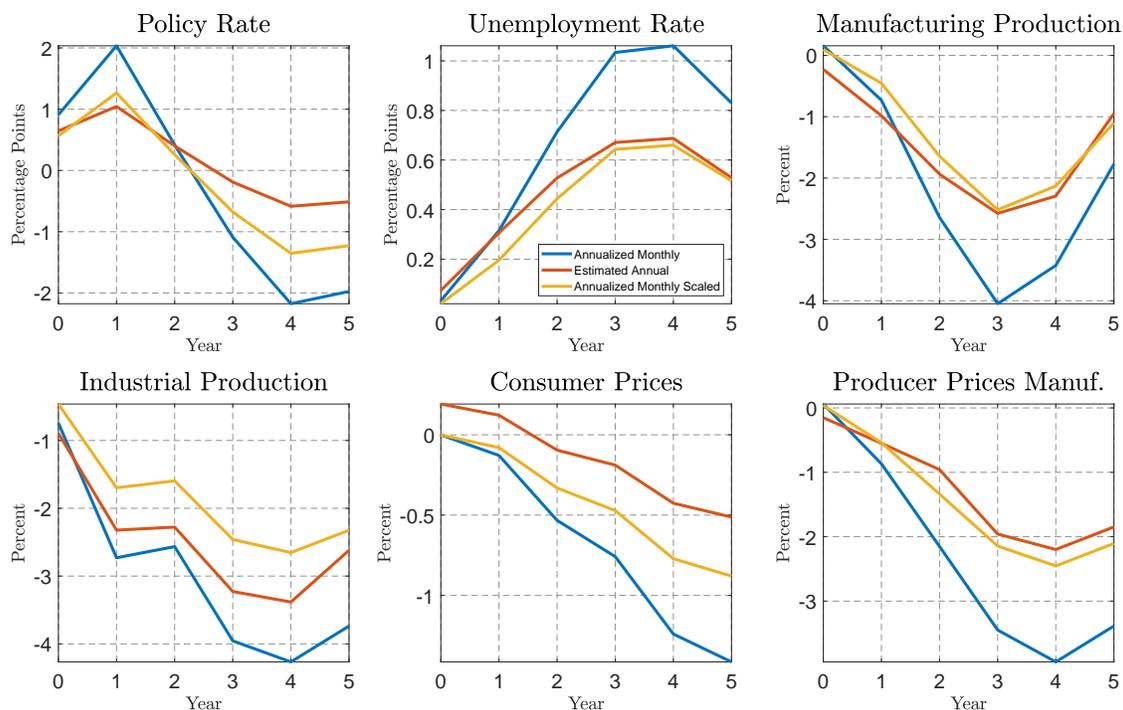


Figure 30: **Time Aggregation.**

Notes: Impulse responses at an annual frequency; see description in text for the computation of the three types of responses.

Next, we rescale the “annualized monthly” responses. The scaling factor is obtained by comparing “estimated annual” and “annualized monthly” of the average policy rate response over the years 0 to 2, the years for which the response is positive. Based on the “estimated annual” responses, the policy rate increases on average by around 70 basis points over this period, and by around 112 basis points based on the “annualized monthly” responses. The scaling factor is given by the ratio of the two, that is, around 0.62. Using this scaling factor, we rescale all “annualized monthly” responses, which are shown in yellow as “annualized monthly scaled” responses. Based on this rescaling, the “estimated annual” and the “annualized monthly scaled” responses for the unemployment rate, manufacturing production, and producer prices (manufacturing) are close to each other. Somewhat larger gaps remain for industrial production and consumer prices. Overall, this exercise shows that the smaller size responses at the annual frequency are largely driven by an attenuation across all variables. In relative terms, the responses are in fact similar, though we also show that this finding does not apply exactly for all variables.

A.8 Monetary Tightenings vs. Easings

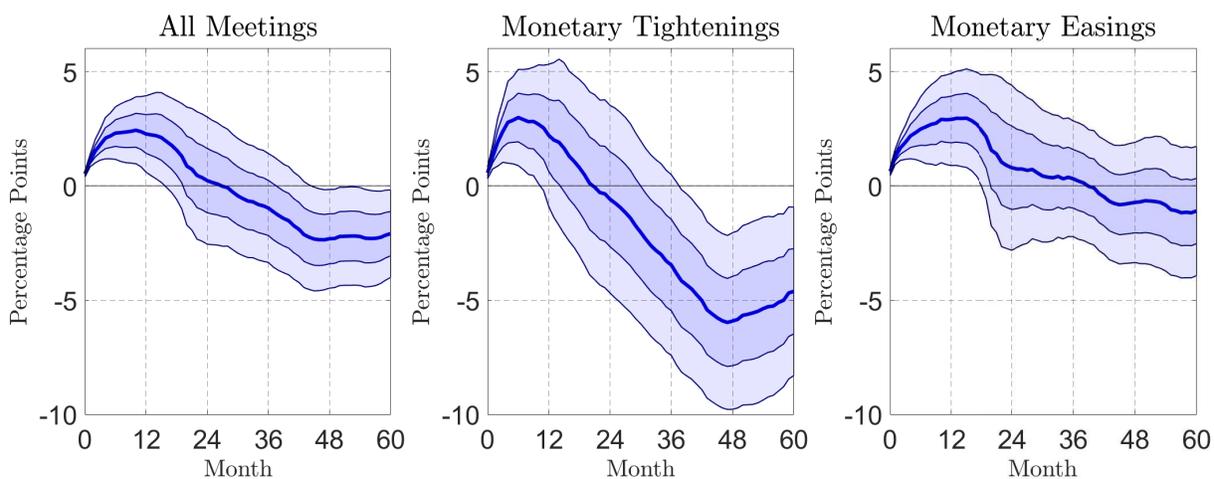


Figure 31: **Policy Rate Responses to Monetary Tightenings & Easings.**

Notes: Impulse responses of the policy rate to a 1 percentage point contractionary monetary policy shock at a monthly frequency, based on the local projection approach in (2). Only positive shocks are included in the series of monetary tightenings, only negative ones in the series of monetary easings. Sample: 1994:M1-2018:M12. 95 and 68 percent confidence bands are shown, using [Newey and West \(1987\)](#) standard errors.

A.9 Comparison with U.S. Data

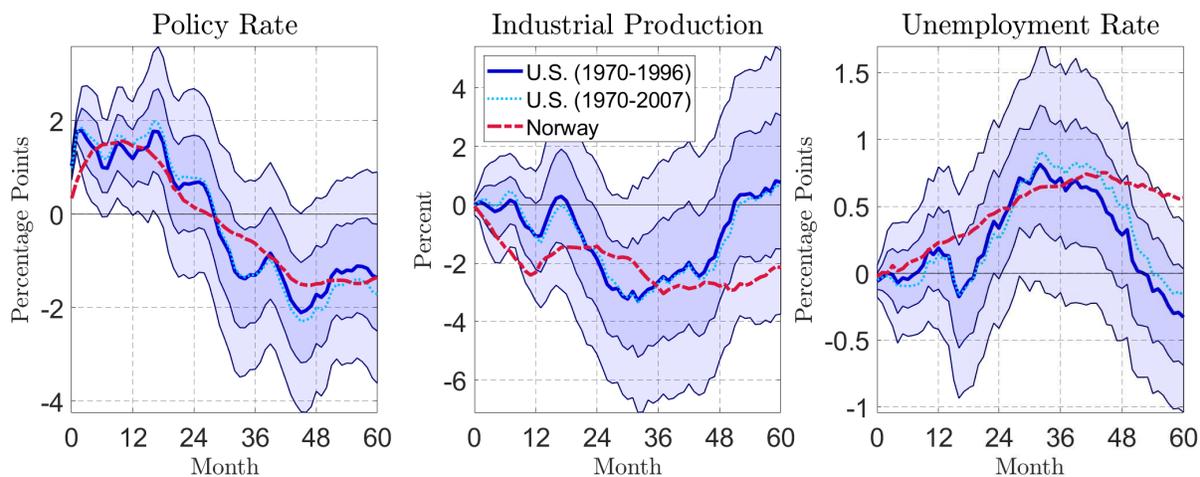


Figure 32: **Comparison of Impulse Responses with U.S. Data.**

Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock taken from Romer and Romer (2004) (1970:M1-1996:M12) and Coibion et al. (2017) (1970:M1-2007:M12), based on the local projection approach in (2). Following Coibion (2012), 1980:M4-M6 and 1980:M9-M11 are excluded. Matching the assumptions based on Norwegian data, three years of lagged values of the monetary policy shocks are added as controls apart from the policy rate. The impulse responses for Norway are the ones shown in Figure (2), but scaled such that the policy rate response matches the one for the United States (1970-1996) after 10 months. 95 and 68 percent confidence bands are shown for the responses based on U.S. data (1970-1996), using Newey and West (1987) standard errors.

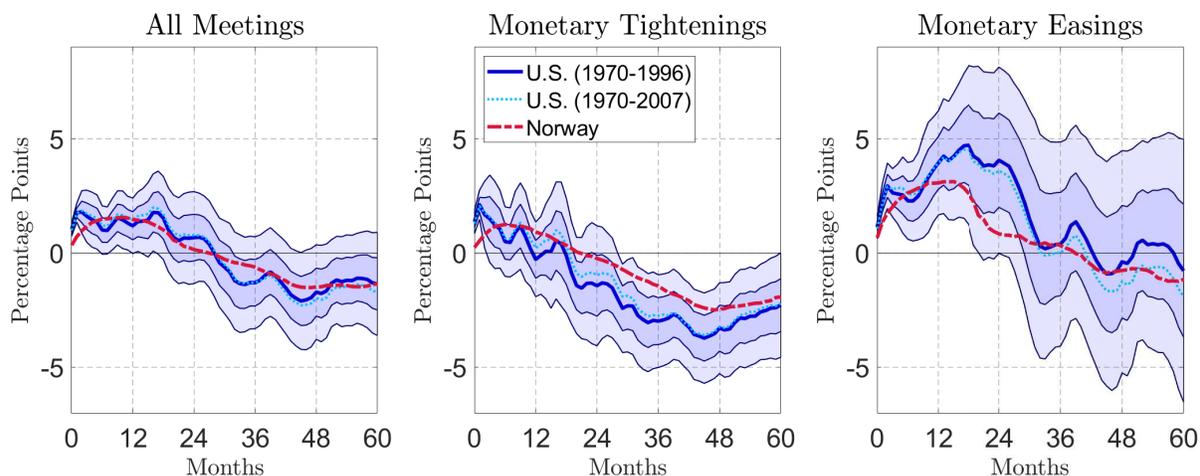


Figure 33: Policy Rate Responses to Monetary Tightenings & Easings – Comparison with U.S. Data.

Notes: Impulse responses of the policy rates to a one percentage point contractionary monetary policy shock taken from [Romer and Romer \(2004\)](#) (1970:M1-1996:M12) and [Coibion et al. \(2017\)](#) (1970:M1-2007:M12), based on local projection approach in (2). Only positive shocks are included in the series of monetary tightenings, only negative ones in the series of monetary easings. Following [Coibion \(2012\)](#), 1980:M4-M6 and 1980:M9-M11 are excluded. The impulse responses for Norway are the ones shown in [Figure \(2\)](#), but scaled such that each respective policy rate response matches the one for the United States (1970-1996) after 10 months. 95 and 68 percent confidence bands are shown for the responses based on U.S. data (1970-1996), using [Newey and West \(1987\)](#) standard errors.

A.10 Robustness

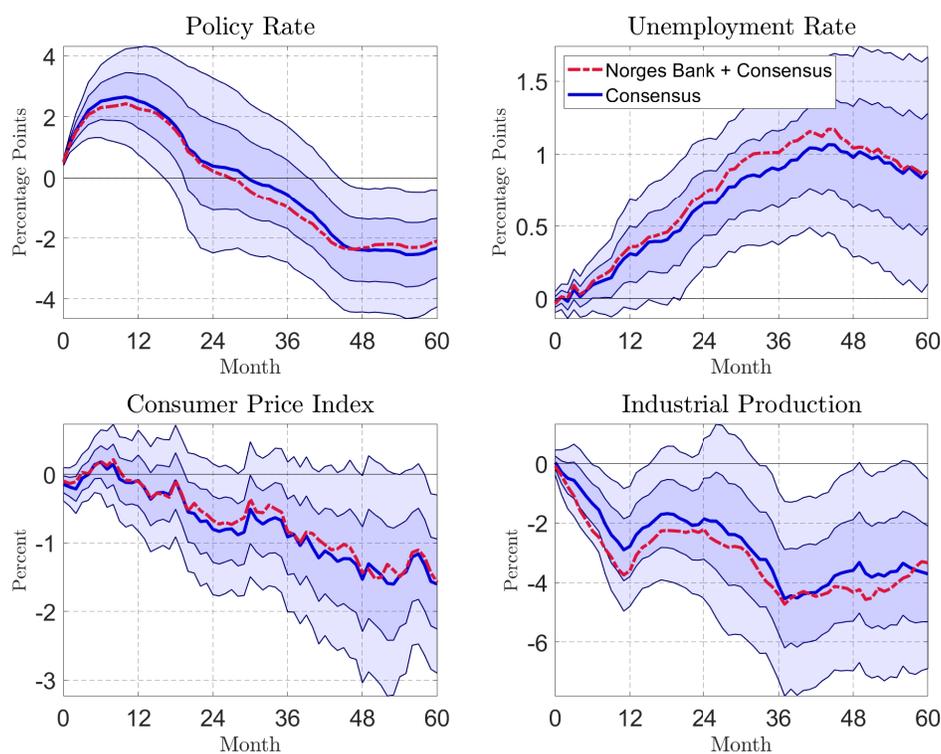


Figure 34: **Impulse Responses – Consensus Forecasts.**

Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at a monthly frequency, based on the local projection approach in (2). “Norges Bank + Consensus” indicates the identification approach used in the main analysis, “Consensus” excludes Norges Bank forecasts leaving the rest unchanged. 95 and 68 percent confidence bands shown, using [Newey and West \(1987\)](#) standard errors. [Table 4](#) in [Appendix A.5](#) lists the data sources.

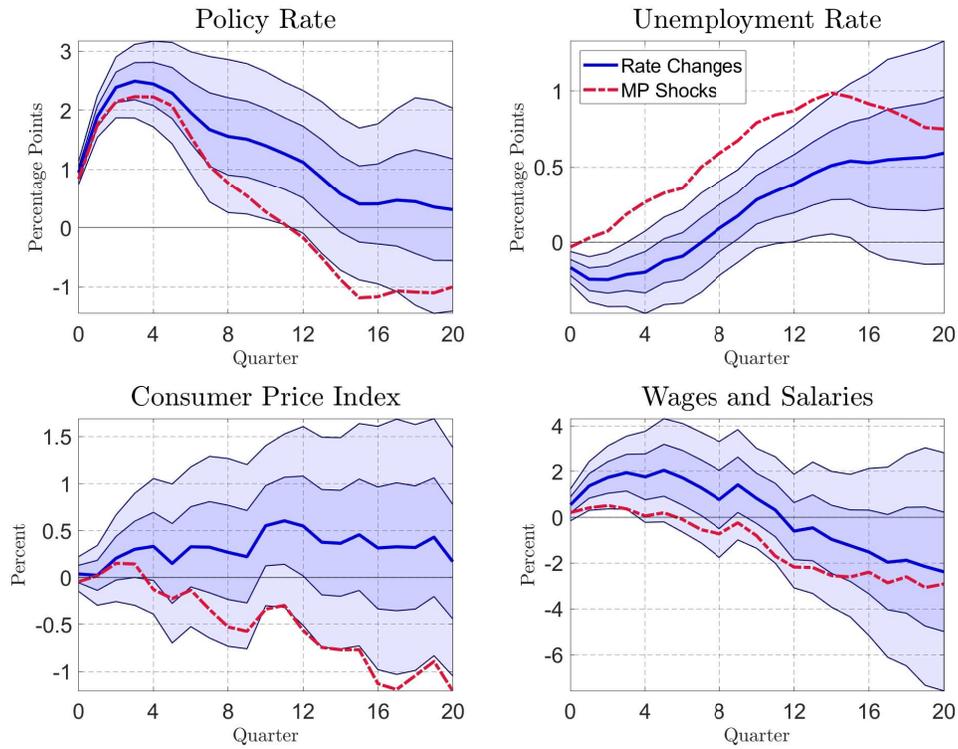


Figure 35: Impulse Responses – Rate Changes.

Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock or rate change at a quarterly frequency, based on the local projection approach in (2). “MP Shocks” indicates the identification approach used in the main analysis, “Rate Changes” uses rate changes instead of shocks in (2). 95 and 68 percent confidence bands shown, using [Newey and West \(1987\)](#) standard errors. Table 4 in Appendix A.5 lists the data sources.

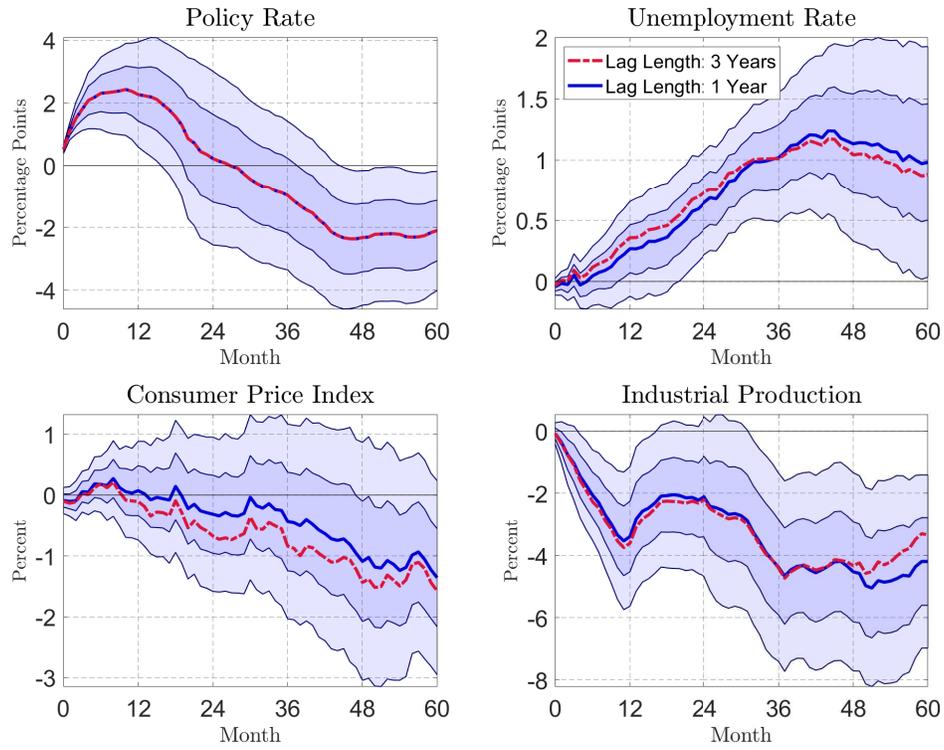


Figure 36: Impulse Responses – Lag Length.

Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at a monthly frequency, based on the local projection approach in (2). “Lag Length: 3 Years” indicates the baseline assumption of three years of lagged shocks as controls, “Lag Length: 1 Year” uses only the shocks from the prior year as controls. 95 and 68 percent confidence bands shown, using [Newey and West \(1987\)](#) standard errors. [Table 4](#) in [Appendix A.5](#) lists the data sources.

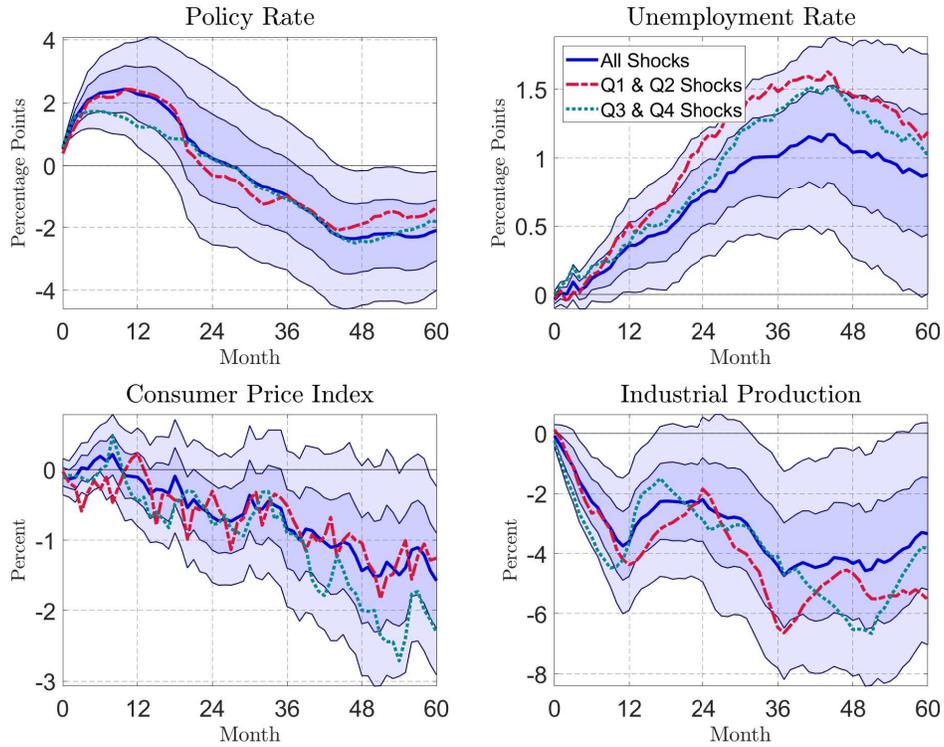


Figure 37: Impulse Responses – Timing of Shocks.

Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at a monthly frequency, based on the local projection approach in (2). “All Shocks” indicates the baseline assumption that considers all monetary policy shocks, “Q1 & Q2 Shocks” excludes shocks that occurred in the second half of a year, and “Q3 & Q4 Shocks” excludes shocks that occurred in the first half of a year. 95 and 68 percent confidence bands shown, using [Newey and West \(1987\)](#) standard errors. Table 4 in Appendix A.5 lists the data sources.

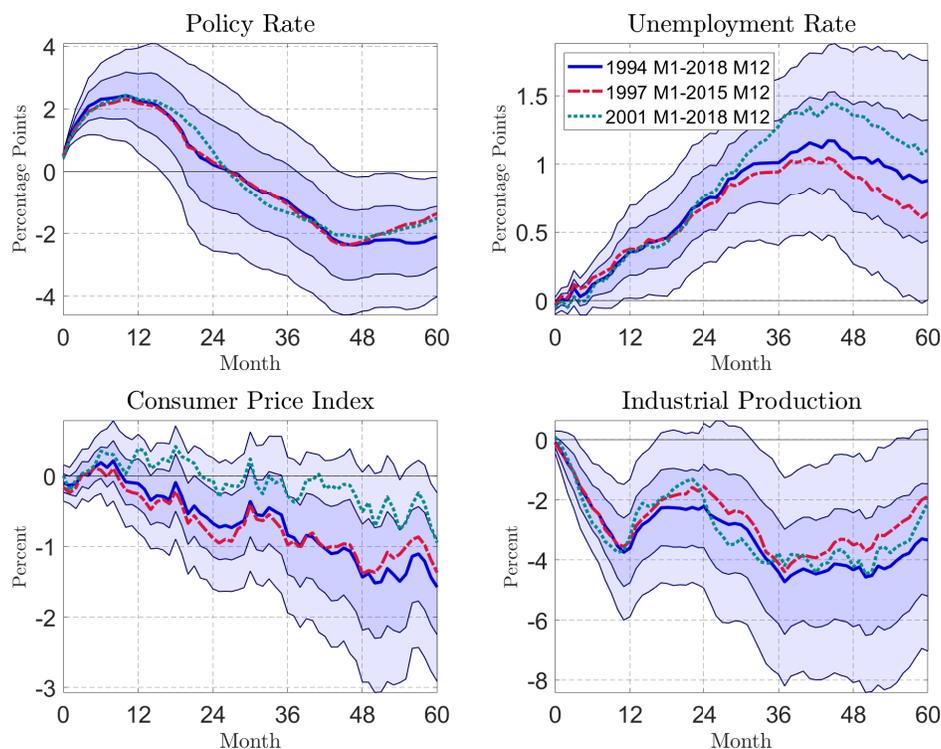


Figure 38: **Impulse Responses – Sample.**

Notes: Impulse responses to a 1 percentage point contractionary monetary policy shock at a monthly frequency, based on the local projection approach in (2). “1994:M1-2018:M12” refers to baseline sample, which are restricted to “1997:M1-2015:M12” and “2001:M1-2018:M12” for the other estimations. 95 and 68 percent confidence bands shown, using [Newey and West \(1987\)](#) standard errors. Table 4 in Appendix A.5 lists the data sources.

B Micro Data

B.1 Details on Consumption Imputation

In this appendix, we describe how we compute capital gains at the household level for four assets: housing, stocks, stock funds, and private businesses.

Housing. For housing, we observe housing transactions directly in the data. Capital gains in housing are the change in housing wealth that is not due to housing transactions.

Stocks. For stocks listed at Oslo Stock Exchange we use two approaches to compute capital gains: (i) when we observe individual level stock ownership (2005-2015) and (ii) when we do not observe individual level stock ownership (1993-2004). For the latter period, we observe each household’s ownership of individual stocks (both volume and price) at the end of each year. In addition, we observe the daily price and turnover of each individual stock from the stock exchange. To compute a measure of capital gains, we make two simplifying assumptions:

1. All transactions are in the same direction throughout the year. That is, if the household is a net seller, it is a net seller throughout the year.
2. A household's share of transactions in a given day is equal to that day's share of transactions as a share of total transactions over the year.

With these two assumptions at hand, we can compute a measure of capital gains for each stock at the household level using only the data available to us. To see this, note that the value of net transactions over the year is equal to $\sum_{n=1}^N p_{j,n,t} dq_{j,n,t}$, where N is the number of days, $p_{j,n,t}$ is the price of stock j on day n in year t , and $dq_{j,n,t}$ is the transactions of the stock on day n . The two assumptions above imply that $dq_{j,n,t} = (q_{j,t} - q_{j,t-1})\omega_{j,n,t}$ where $q_{j,t}$ is the number of shares at the end of the year, $q_{j,t-1}$ is the number of shares at the beginning of the year, and $\omega_{j,n,t}$ is the share of yearly transactions taking place on day n for stock j in year t . We can therefore compute a measure of net transactions as

$$\sum_{n=1}^N p_{j,n,t} dq_{j,n,t} = \sum_{n=1}^N p_{j,n,t} (q_{j,t} - q_{j,t-1}) \omega_{j,n,t} = (q_{j,t} - q_{j,t-1}) \sum_{n=1}^N p_{j,n,t} \omega_{j,n,t}. \quad (8)$$

Capital gains for an individual asset are then equal to $p_{j,t}q_{j,t} - p_{j,t-1}q_{j,t-1} - (q_{j,t} - q_{j,t-1}) \sum_{n=1}^N p_{j,n,t} \omega_{j,n,t}$ and can be obtained from observables in our data.

For the early period, 1993-2005, for which we do not observe individual stock ownership, we utilize the same method as above but apply it on the total holdings of stocks. Between 1993 and 2005, the capital gains for stocks in year t are equal to

$$p_t q_t - p_{t-1} q_{t-1} - (q_t - q_{t-1}) \sum_{n=1}^N p_{n,t} \omega_{n,t} = p_t q_t - p_{t-1} q_{t-1} - (p_t q_t - \frac{p_t}{p_{t-1}} p_{t-1} q_{t-1}) \sum_{n=1}^N \frac{p_{n,t}}{p_t} \omega_{n,t} \quad (9)$$

where p_t is the price of the stock index at the end of the year, q_t is the number of shares at the end of the year, $p_{n,t}$ is the price on day n in year t , and $\omega_{n,t}$ is the share of yearly transactions on day n in year t . Note that we do not observe the number of shares for stocks in total, only total wealth in stocks ($p_t q_t$) and the price movements (p_t / p_{t-1}). We are therefore using the expression for capital gains after the equality sign in (9).

Stock Funds. We do not have access to household-level holdings of individual stock funds. We therefore use a measure of capital gains on stock funds from the Financial Accounts from 1996 to 2015. In particular, we assume that the capital gains on the stock funds portfolio of each household obtain the same percentage change as the total from the Financial Accounts.

Private Business. We attribute no capital gains to private businesses unless they hold stocks on the Oslo Stock Exchange. For the case that a private business owns stocks, these stocks are allocated to household portfolios using the ownership registry of private businesses before computing capital gains on those stocks using the method described above.

B.2 Details on Institutional Setting

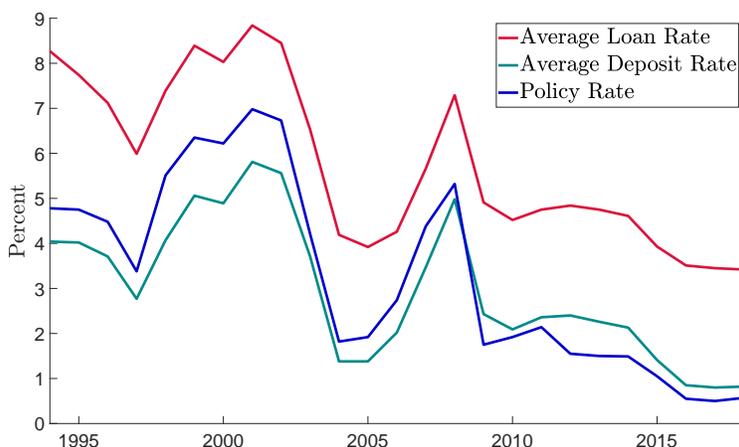


Figure 39: **Historical Average Policy Rate, Loan Rate, and Deposit Rate.**

Notes: This figure shows the weighted average interest rates on all outstanding loans and deposits in Norwegian banks together with central bank policy rate. Table 4 in Appendix A.5 lists the data sources.

C Monetary Policy Transmission at the Household Level

C.1 Inequality

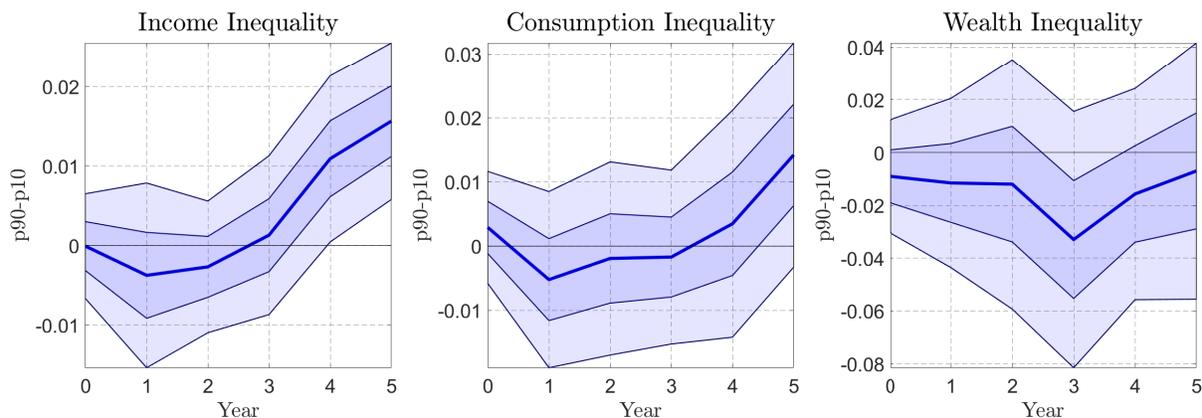


Figure 40: **Impulse Responses of Inequality Measures.**

Notes: Impulse responses of p90-p10 measure to a 1 percentage point contractionary monetary policy shock at an annual frequency, based on the local projection approach in (2). The p90-p10 measure is the difference between the 90th and 10th percentile of log levels of each distribution. 95 and 68 percent confidence bands shown, using Newey and West (1987) standard errors.

C.2 Baseline Results

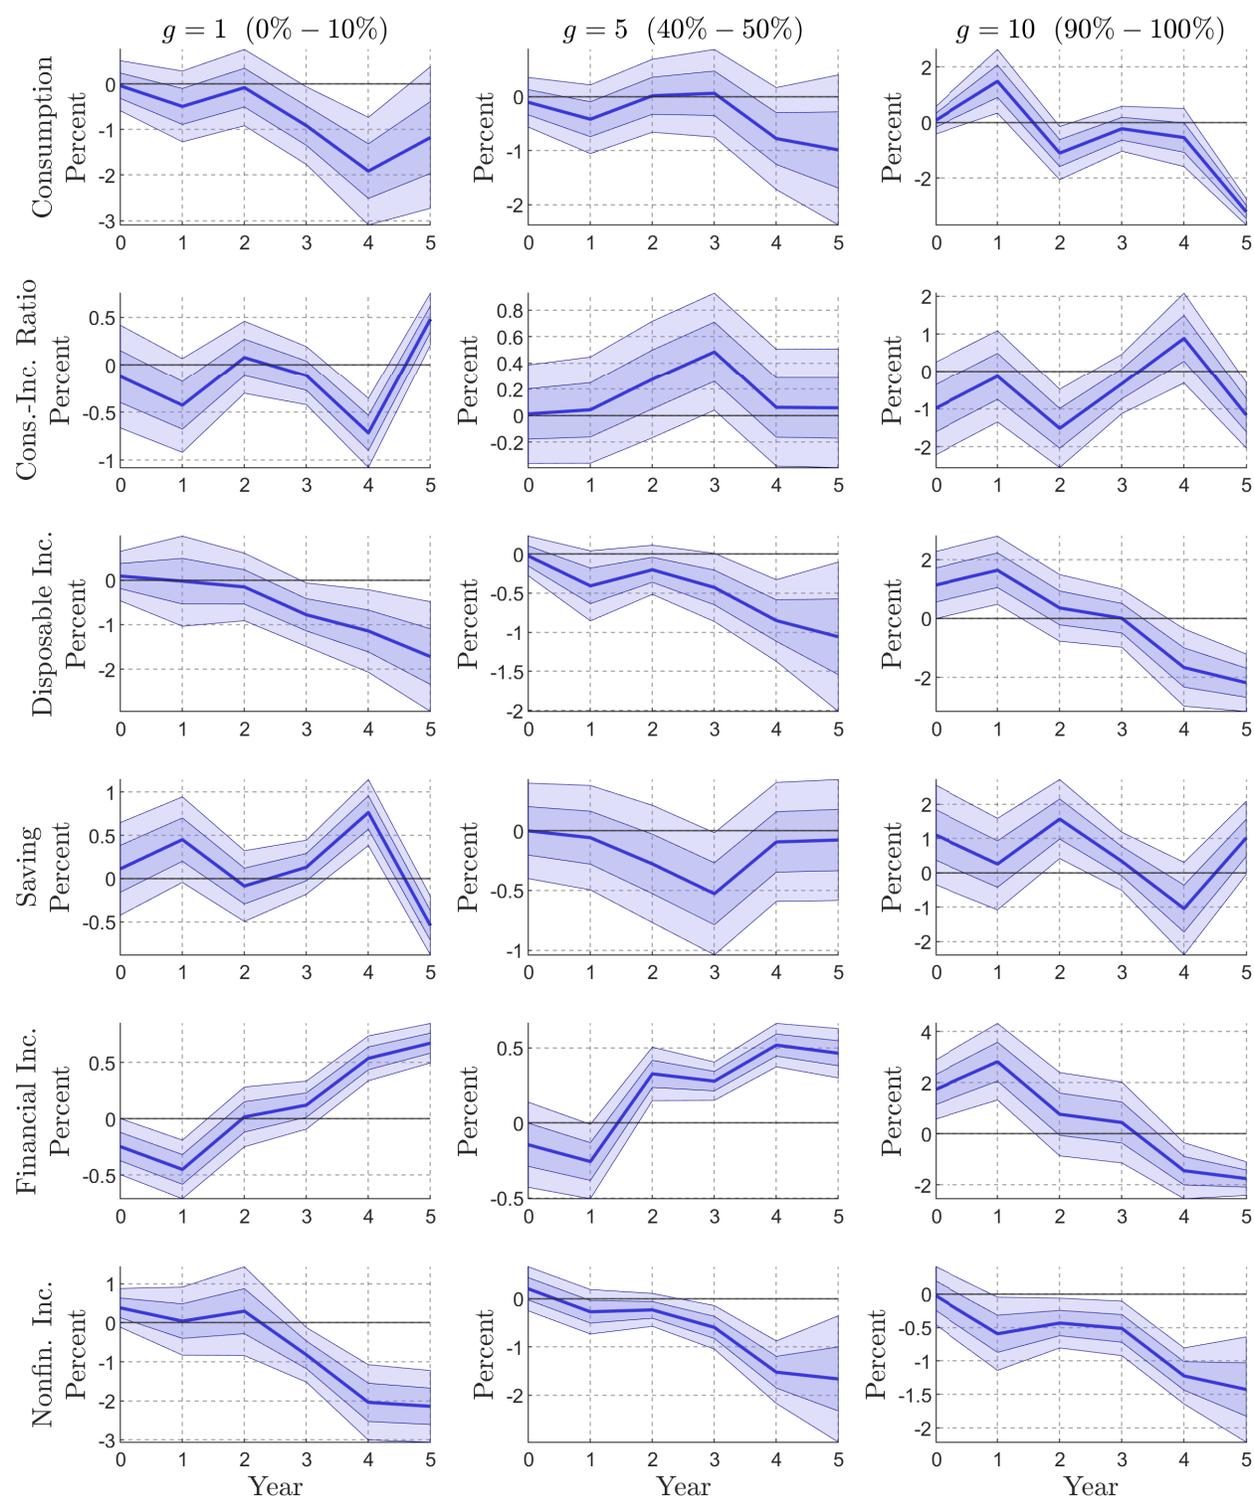


Figure 41: Impulse Responses of Selected Groups.

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. 95 and 68 percent confidence bands shown, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

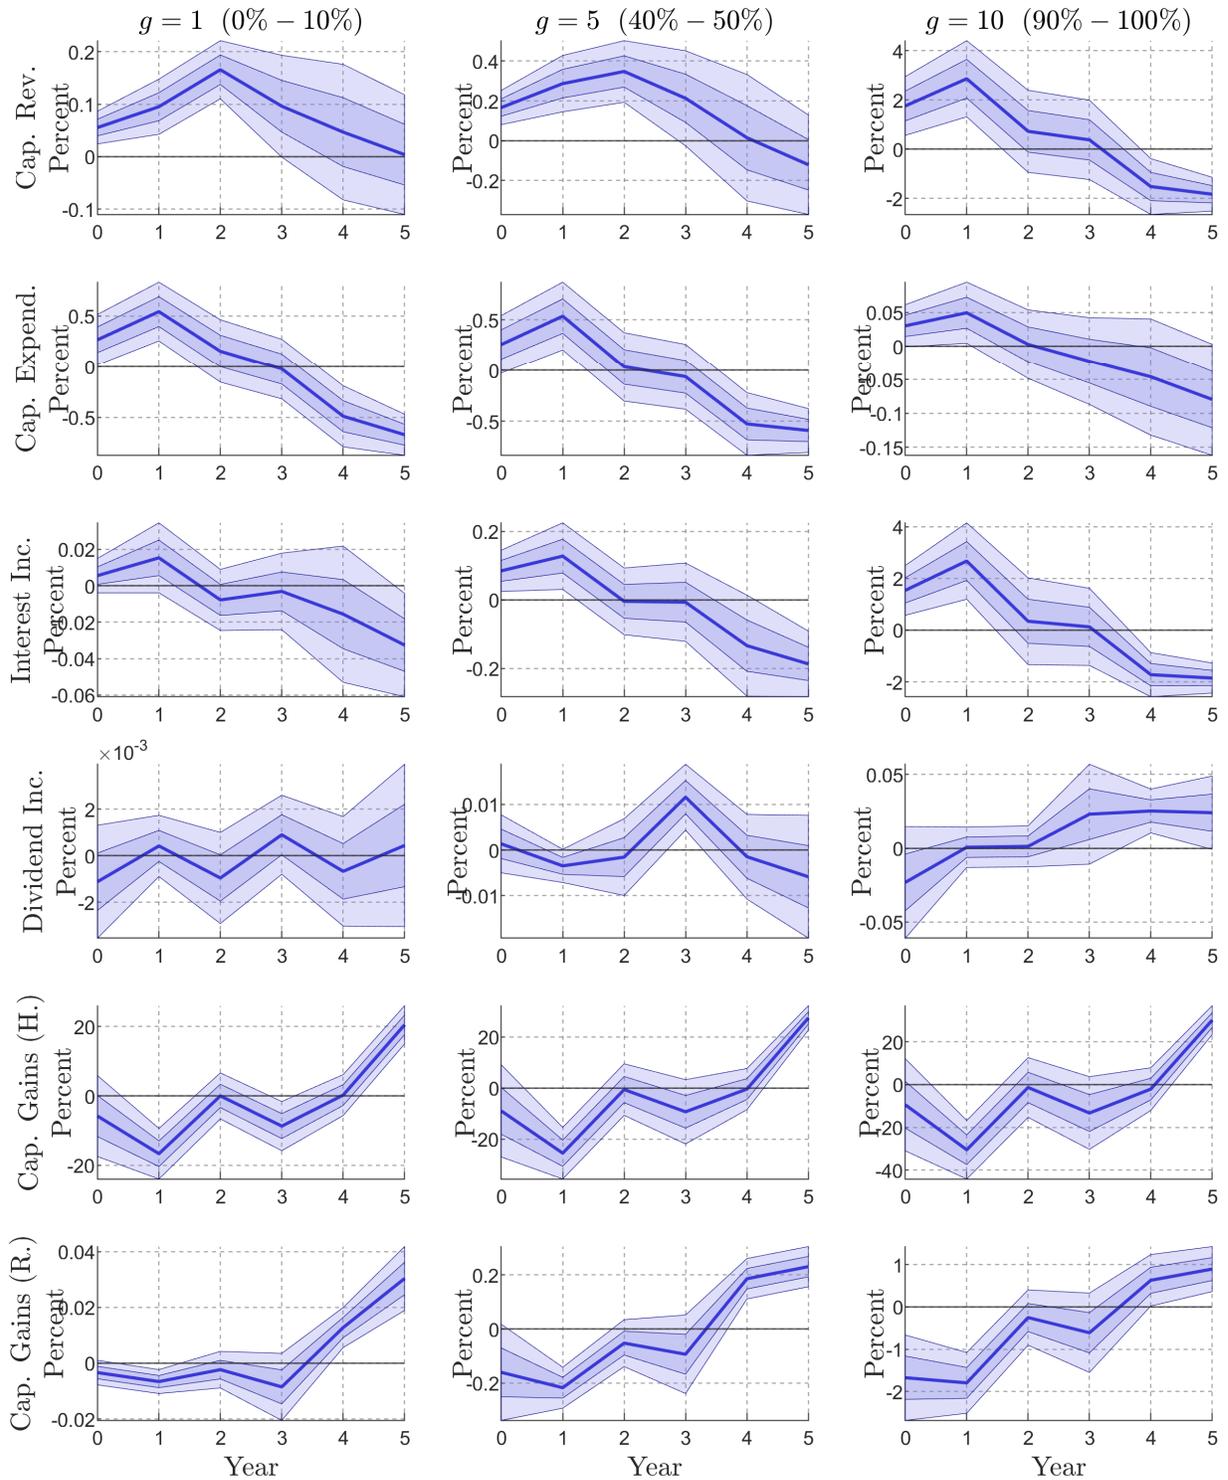


Figure 42: Impulse Responses of Selected Groups (Continued).

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. 95 and 68 percent confidence bands shown, computed using Driscoll and Kraay (1998) standard errors.

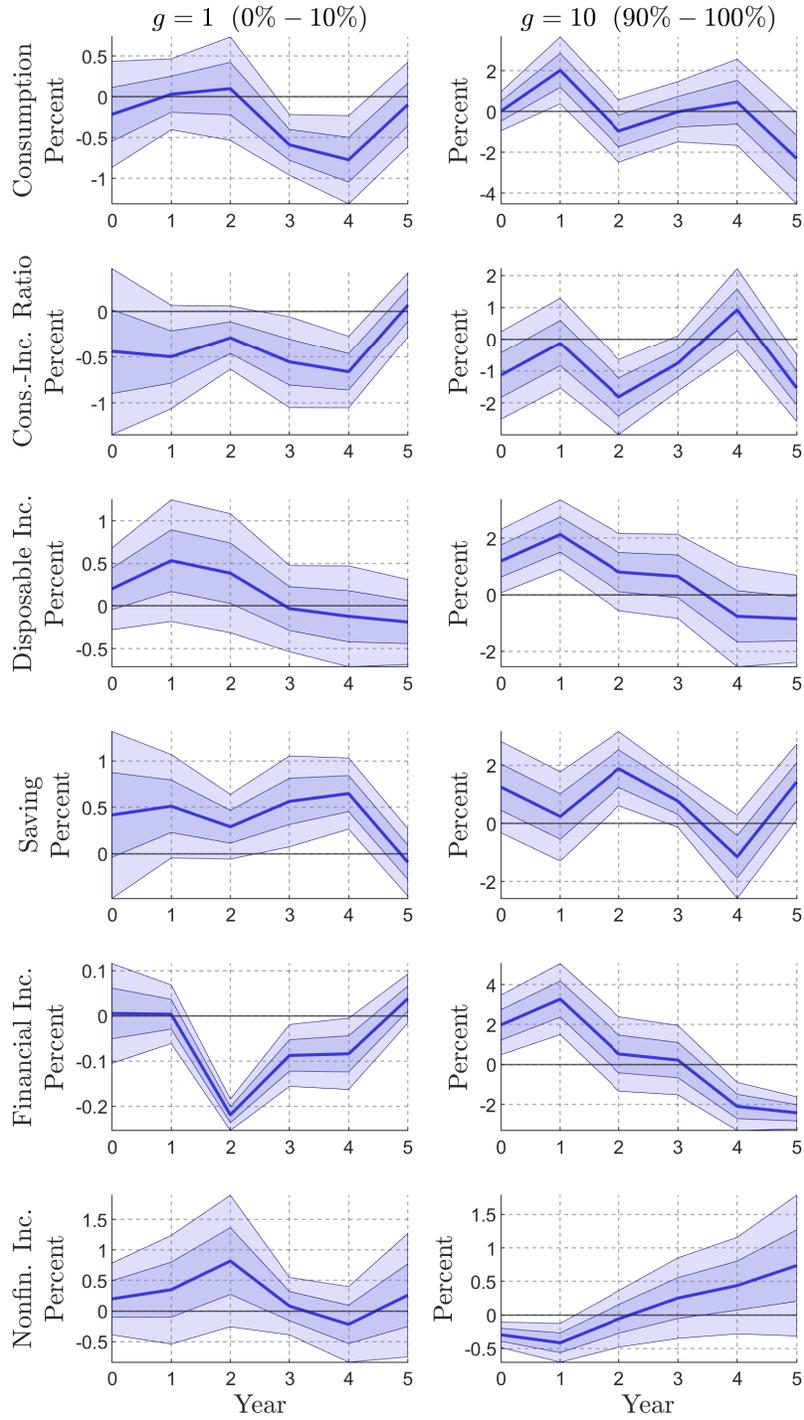


Figure 43: **Impulse Responses at Bottom and Top Relative to Group at Median.**

Notes: Differences between impulse responses of 0-10 percent group and 40-50 percent group (left column); differences between impulse responses of 90-100 percent group and 40-50 percent group (right column). Changes relative to lagged net income at segments of the liquid asset distribution.

The model estimated is given by $(y_{i,t+h} - y_{i,t-1})/inc_{i,t-1} = \delta_{i,g}^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{g=1,10} \beta_g^h D_{i,t-1}^g \cdot \epsilon_t^{MP} + \sum_{g=1,5,10} \sum_{k=1}^K \gamma_{g,k}^h X_{i,t-k} D_{i,t-1}^g + u_{i,t}^h$ for all i in groups $g = 1, 5, 10$, where $D_{i,t-1}^g$ equals 1 if i is in g in $t-1$. β_1^h and β_{10}^h shown with 95 and 68 percent confidence bands based on Driscoll and Kraay (1998) strd. errors.

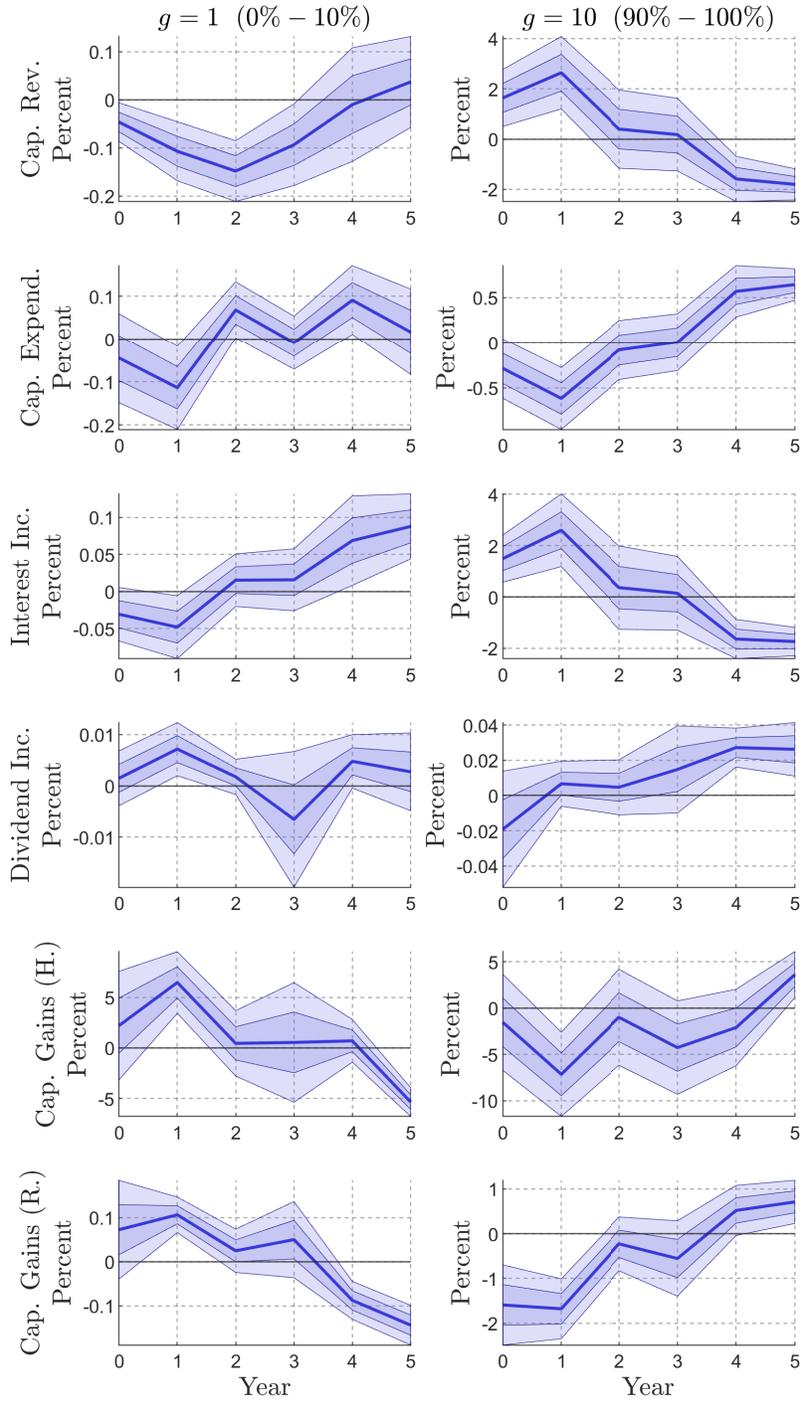


Figure 44: Impulse Responses at Bottom and Top Relative to Group at Median (Continued).

Notes: Differences between impulse responses of 0-10 percent group and 40-50 percent group (left column); differences between impulse responses of 90-100 percent group and 40-50 percent group (right column). Changes relative to lagged net income at segments of the liquid asset distribution.

The model estimated is given by $(y_{i,t+h} - y_{i,t-1})/inc_{i,t-1} = \delta_{i,g}^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{g=1,10} \beta_g^h D_{i,t-1}^g \cdot \epsilon_t^{MP} + \sum_{g=1,5,10} \sum_{k=1}^K \gamma_{g,k}^h X_{i,t-k} D_{i,t-1}^g + u_{i,t}^h$ for all i in groups $g = 1, 5, 10$, where $D_{i,t-1}^g$ equals 1 if i is in g in $t - 1$. β_1^h and β_{10}^h shown with 95 and 68 percent confidence bands based on [Driscoll and Kraay \(1998\)](#) strd. errors.

C.3 Robustness

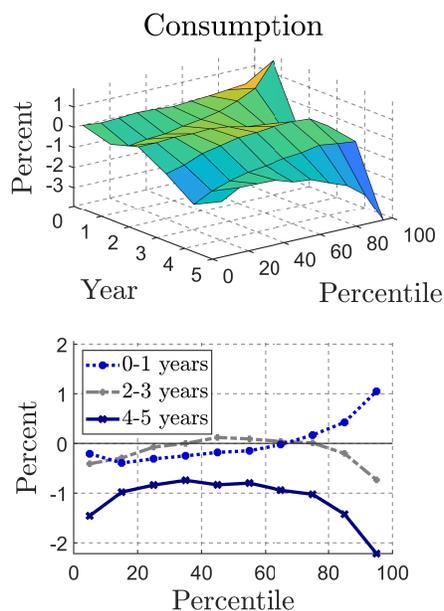


Figure 45: **Impulse Responses of Consumption Growth.**

Notes: Dependent variable given by cumulative growth rate $(c_{i,t+h} - c_{i,t-1}) / c_{i,t-1}$. Model otherwise as specified in equation (5). Plotted are percentage changes at segments of the liquid asset distribution.

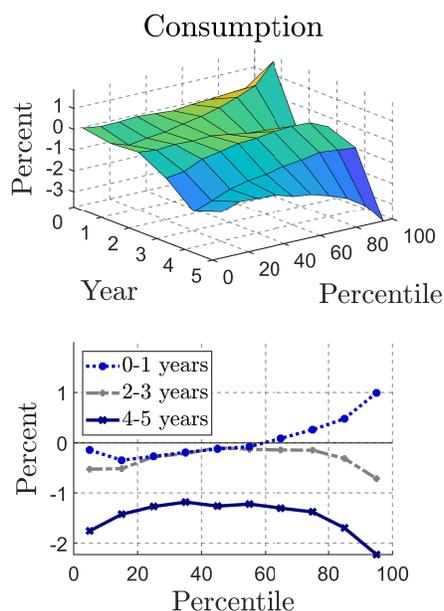


Figure 46: **Impulse Responses of Consumption with Housing Service Flow.**

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. Housing service flow computed as ratio of aggregate housing service flow from owner-occupied housing in national accounts and aggregate housing wealth multiplied by household's beginning-of-year housing wealth. Aggregate housing service flow ratio fell from 3.7 percent in 1995 to 1.8 percent in 2015 (average 2.3 percent).

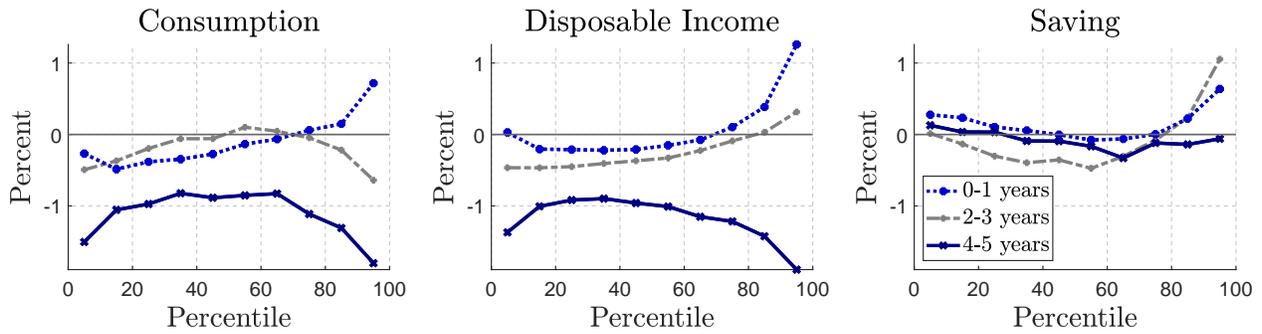


Figure 47: Impulse Responses by Lagged Liquidity-to-Income Ratio.

Notes: Changes relative to lagged net income at segments of the distribution of liquid assets over income.

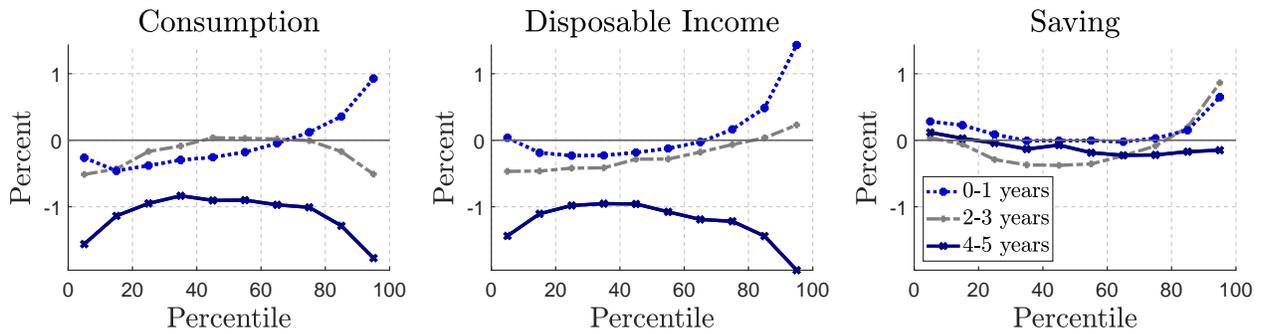


Figure 48: Impulse Responses for Nonstockholders.

Notes: Changes relative to lagged net income at segments of the liquid asset distribution.

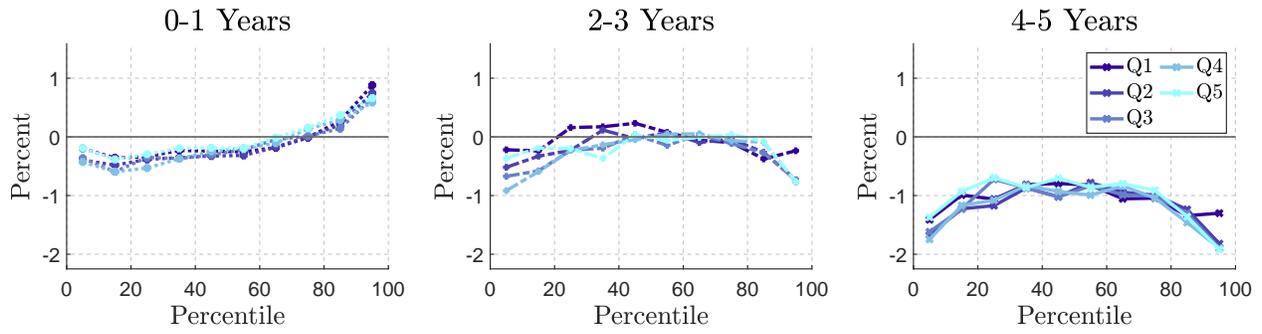


Figure 49: Consumption Responses by Lagged Income Quintile.

Notes: Changes relative to lagged net income at segments of the liquid asset distribution. Q1 to Q5 are quintiles of the income distribution.

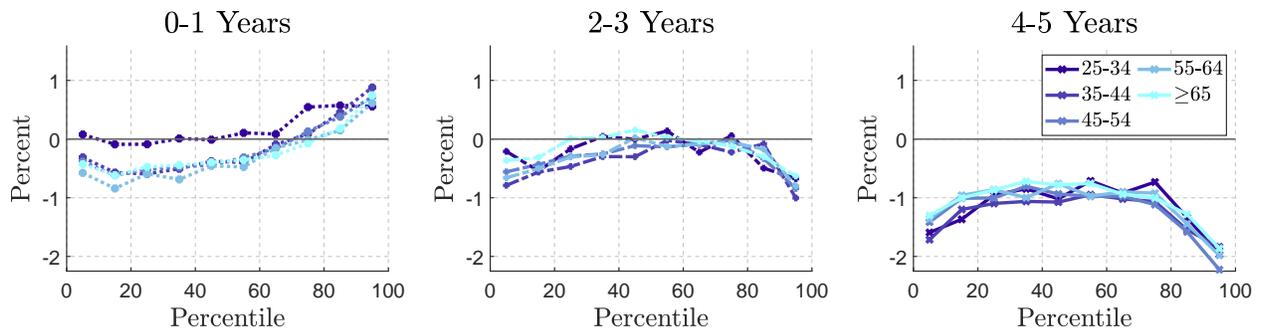


Figure 50: Consumption Responses by Age.

Notes: Changes relative to lagged net income at segments of the liquid asset distribution for five age groups.

C.4 Decomposing the Aggregate Effects of Monetary Policy

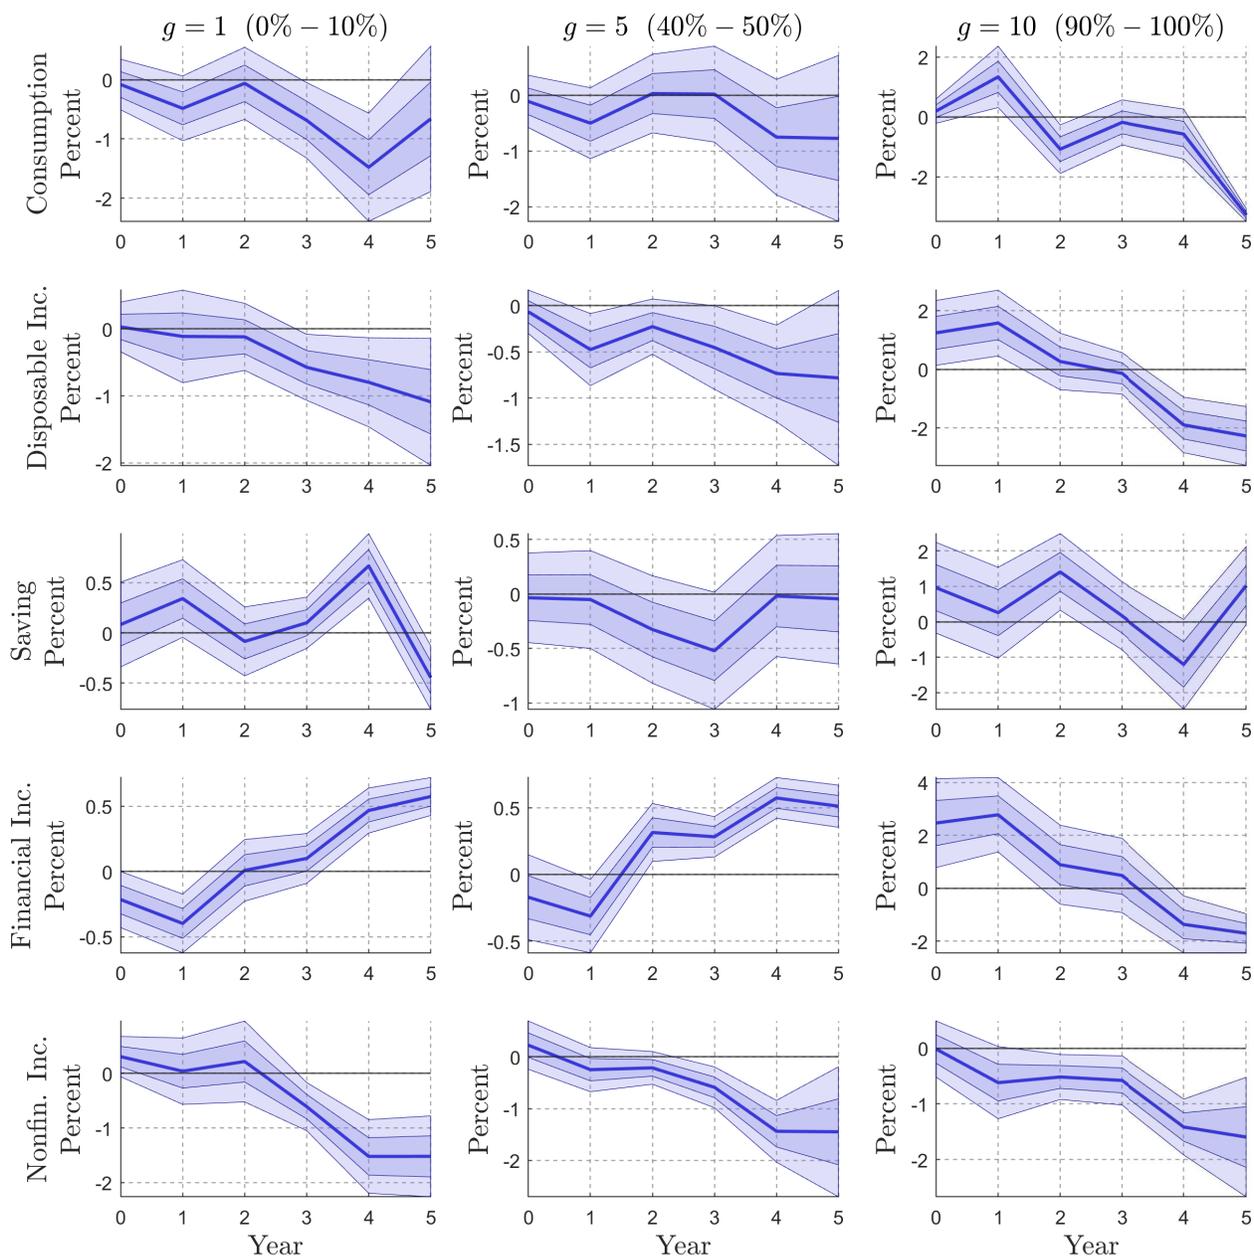


Figure 51: Impulse Responses of Selected Groups in Units of Average Net Income.

Notes: Changes relative to average lagged net income at segments of the liquid asset distribution. 95 and 68 percent confidence bands shown, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

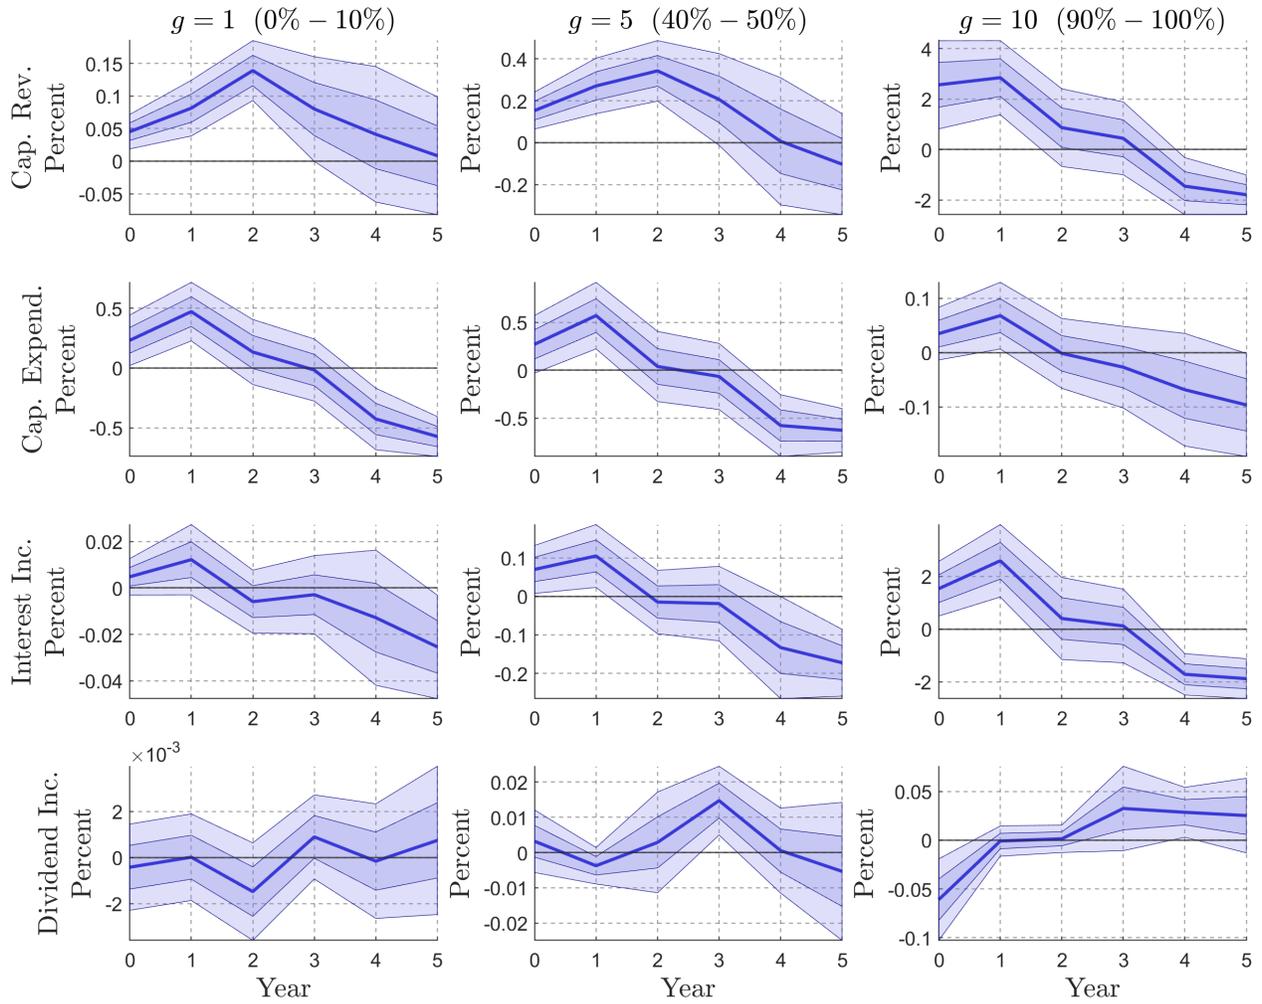


Figure 52: **Impulse Responses of Selected Groups in Units of Average Net Income (Continued).**
Notes: Changes relative to average lagged net income at segments of the liquid asset distribution. 95 and 68 percent confidence bands shown, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

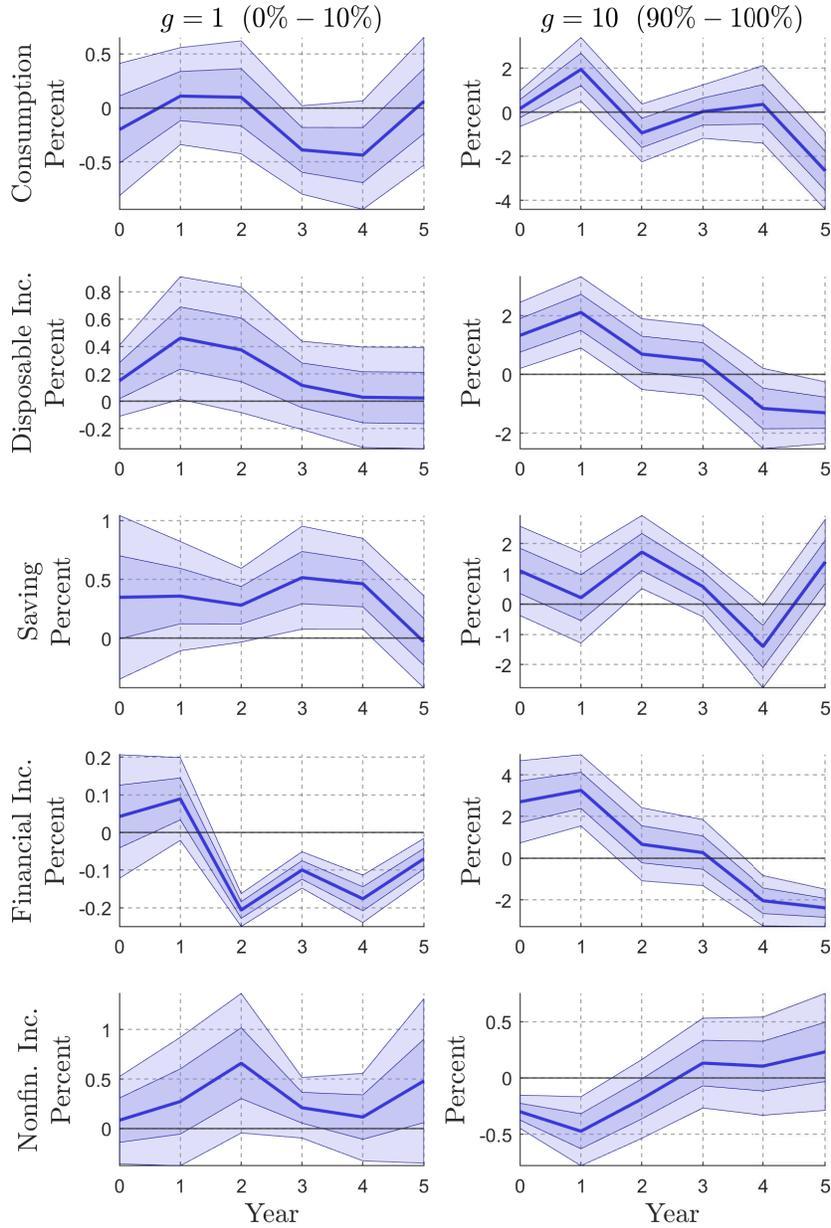


Figure 53: **Impulse Responses at Bottom and Top Relative to Group at Median in Units of Average Net Income.**

Notes: Differences between impulse responses of 0-10 percent group and 40-50 percent group (left column); differences between impulse responses of 90-100 percent group and 40-50 percent group (right column). Changes relative to average lagged net income at segments of the liquid asset distribution.

The model estimated is given by $(y_{i,t+h} - y_{i,t-1})/\overline{inc}_{i,t-1} = \delta_{i,g}^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{g=1,10} \beta_g^h D_{i,t-1}^g \cdot \epsilon_t^{MP} + \sum_{g=1,5,10} \sum_{k=1}^K \gamma_{g,k}^h X_{i,t-k} D_{i,t-1}^g + u_{i,t}^h$ for all i in groups $g = 1, 5, 10$, where $D_{i,t-1}^g$ equals 1 if i is in g in $t - 1$. Shown are β_1^h and β_{10}^h with 95 and 68 percent confidence bands, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

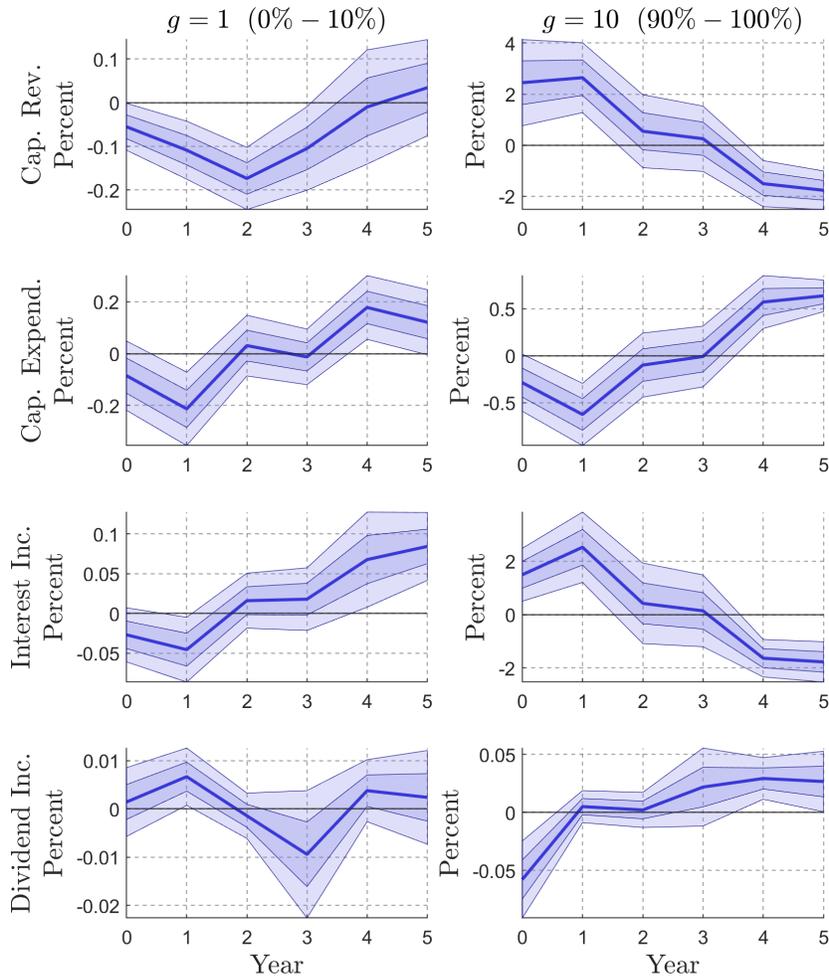


Figure 54: **Impulse Responses at Bottom and Top Relative to Group at Median in Units of Average Net Income (Continued).**

Notes: Differences between impulse responses of 0-10 percent group and 40-50 percent group (left column); differences between impulse responses of 90-100 percent group and 40-50 percent group (right column). Changes relative to average lagged net income at segments of the liquid asset distribution.

The model estimated is given by $(y_{i,t+h} - y_{i,t-1})/\overline{inc}_{i,t-1} = \delta_{i,g}^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{g=1,10} \beta_g^h D_{i,t-1}^g \cdot \epsilon_t^{MP} + \sum_{g=1,5,10} \sum_{k=1}^K \gamma_{g,k}^h X_{i,t-k} D_{i,t-1}^g + u_{i,t}^h$ for all i in groups $g = 1, 5, 10$, where $D_{i,t-1}^g$ equals 1 if i is in g in $t - 1$. Shown are β_1^h and β_{10}^h with 95 and 68 percent confidence bands, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

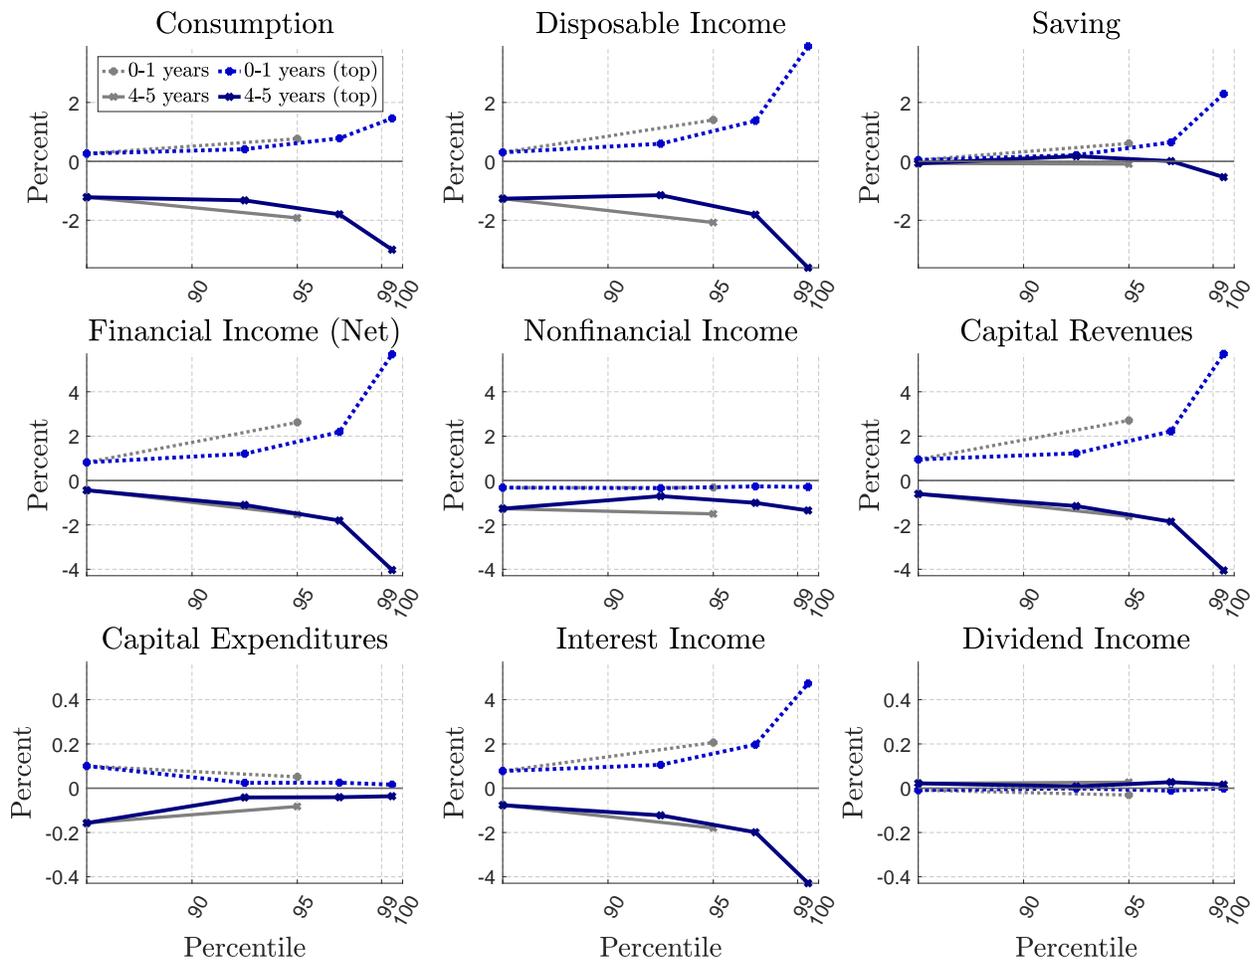


Figure 55: Impulse Responses for Top 10 Percent of Liquid Asset Distribution.

Notes: Changes relative to average lagged net income along the liquid asset distribution. Estimates for $g = 9$ and $g = 10$ against responses of groups 90-95, 95-99, and 99-100 (labeled “top”). All estimates plotted at midpoint of respective segment. See Figures 56-59 for group responses and group comparisons with confidence bands.

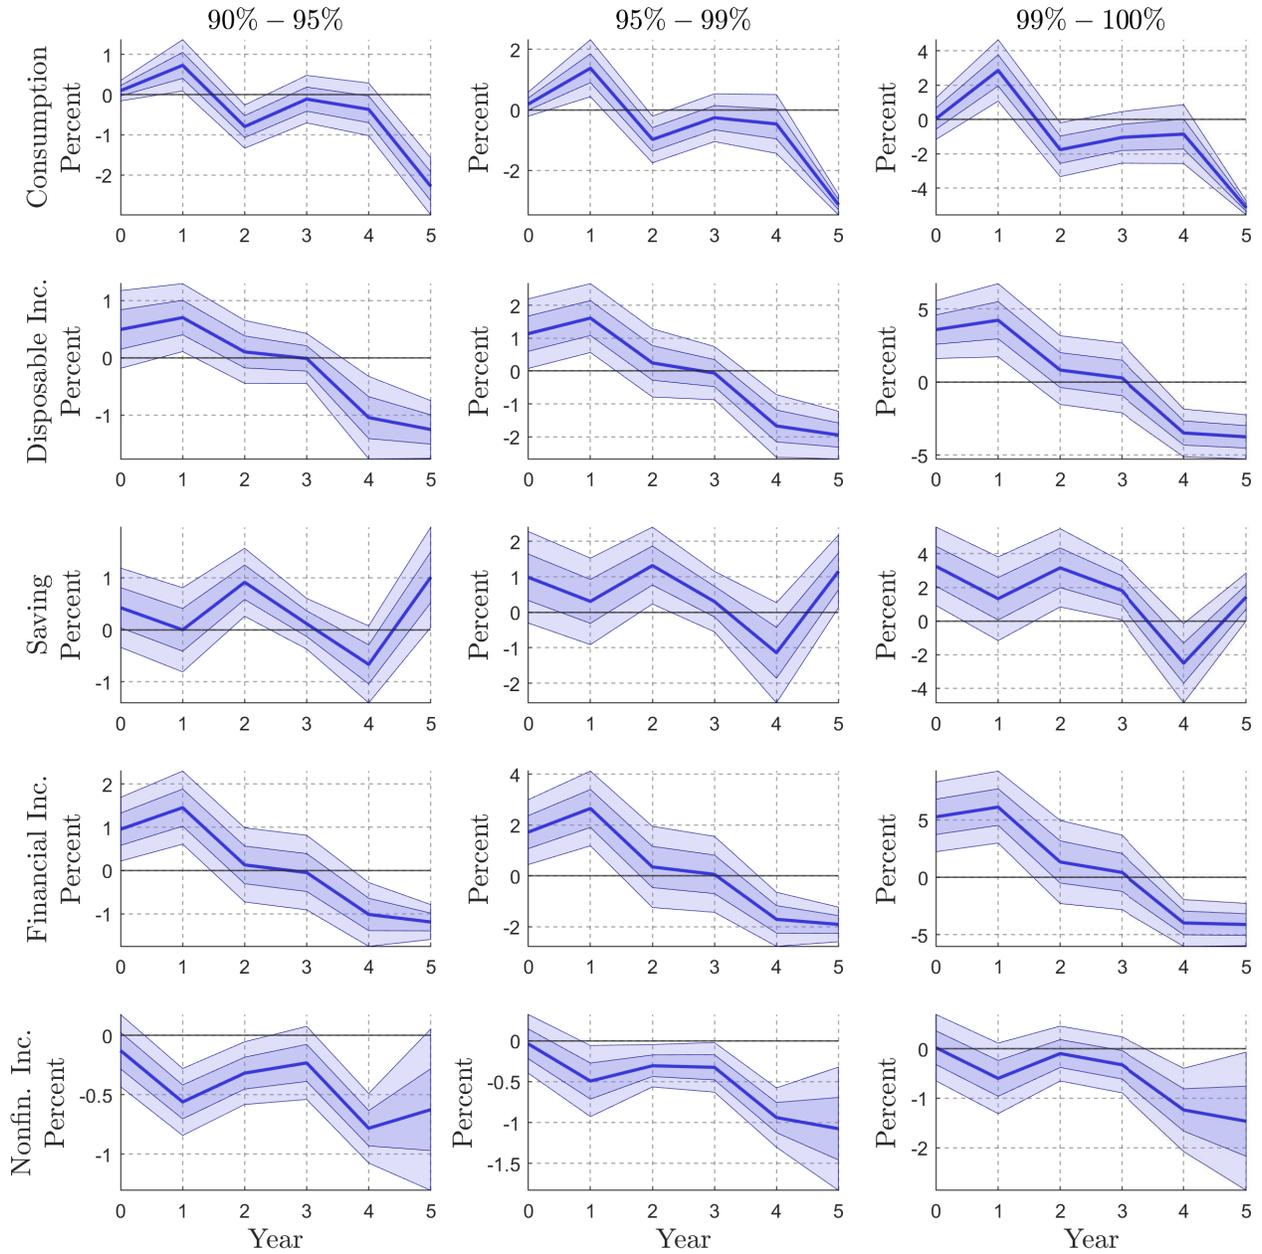


Figure 56: Impulse Responses among Top 10 Percent in Units of Average Net Income.
Notes: Changes relative to average lagged net income at segments of the liquid asset distribution. 95 and 68 percent confidence bands shown, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

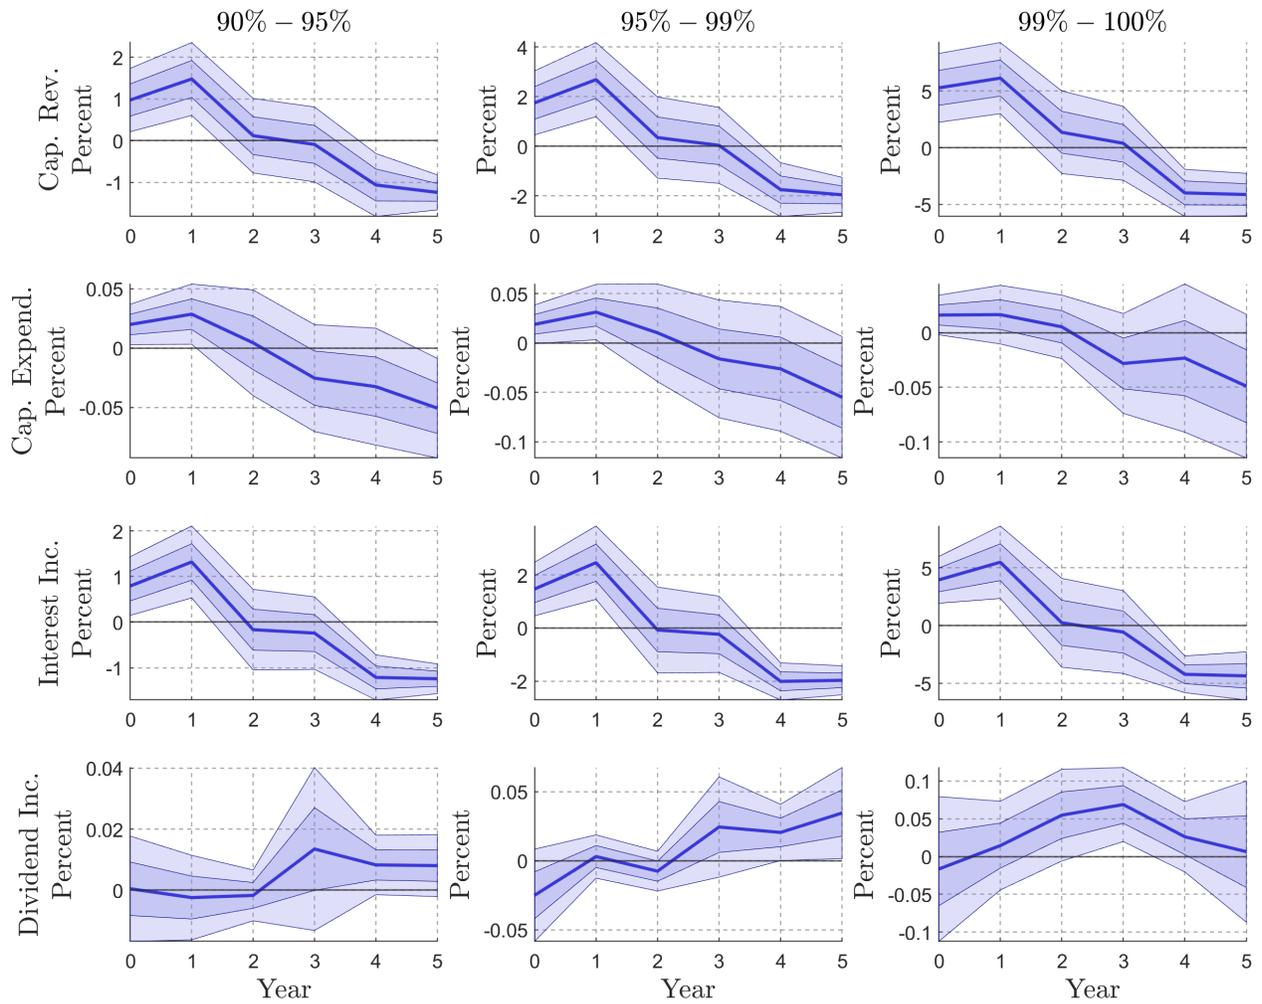


Figure 57: **Impulse Responses among Top 10 Percent in Units of Average Net Income (Continued).**

Notes: Changes relative to average lagged net income at segments of the liquid asset distribution. 95 and 68 percent confidence bands shown, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

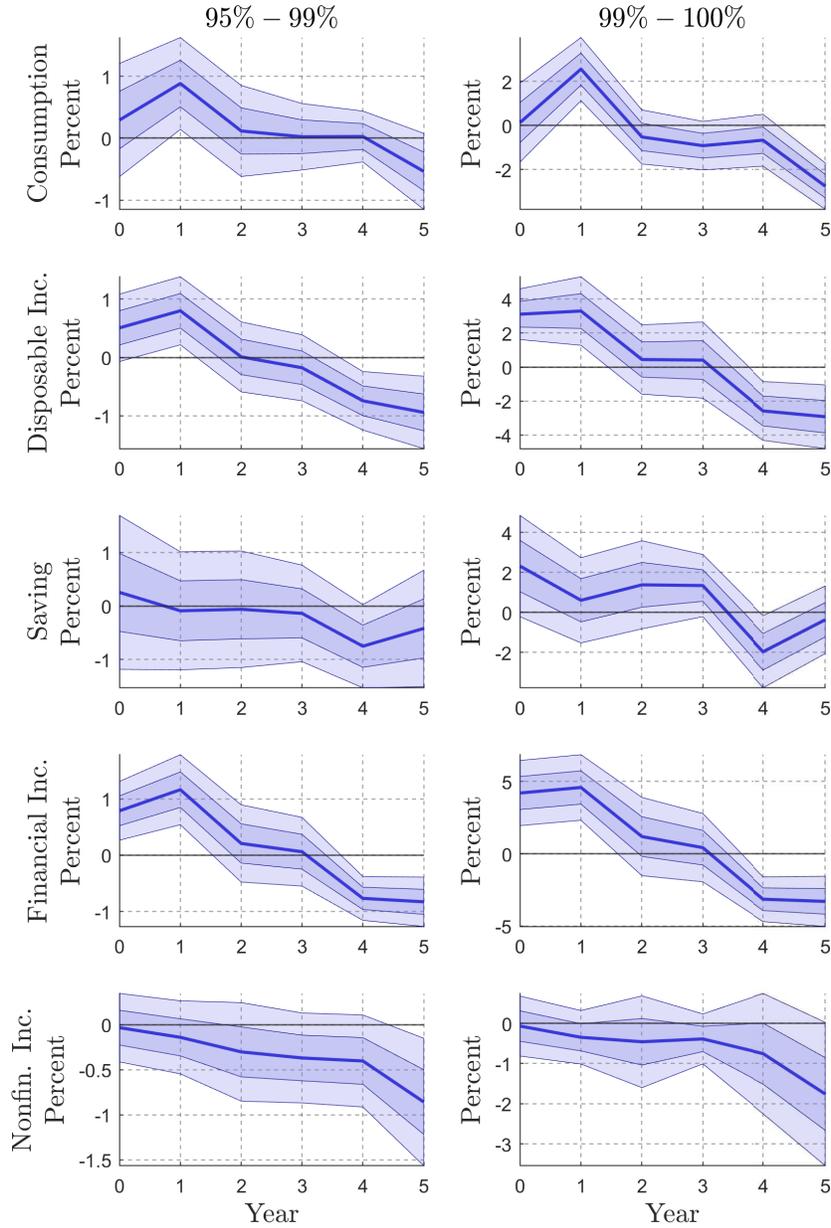


Figure 58: **Impulse Responses at Top Relative to 90-95 Percent Group in Units of Average Net Income.**

Notes: Differences between impulse responses of 95-99 percent group and 90-95 percent group (left column); differences between impulse responses of 99-100 percent group and 90-95 percent group (right column). Changes relative to average lagged net income at segments of the liquid asset distribution.

The model estimated is given by $(y_{i,t+h} - y_{i,t-1})/\bar{inc}_{i,t-1} = \delta_{i,l}^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{l=2,3} \beta_l^h D_{i,t-1}^l \cdot \epsilon_t^{MP} + \sum_{l=1,2,3} \sum_{k=1}^K \gamma_{l,k}^h X_{i,t-k} D_{i,t-1}^l + u_{i,t}^h$ for all i in groups $l = 1, 2, 3$, where $D_{i,t-1}^l$ equals 1 if i is in l in $t - 1$. Groups considered are percentiles 90-95 ($l = 1$), 95-99 ($l = 2$), and 99-100 ($l = 3$). Shown are β_2^h and β_3^h with 95 and 68 percent confidence bands, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

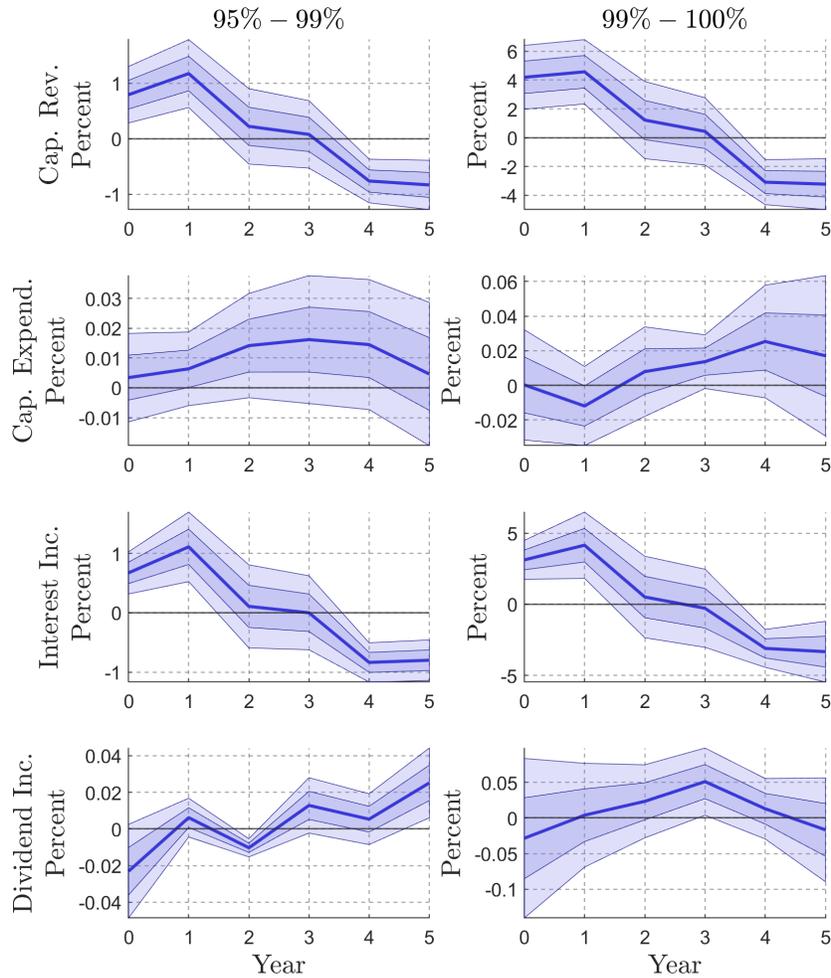


Figure 59: Impulse Responses at Top Relative to 90-95 Percent Group in Units of Average Net Income (Continued).

Notes: Differences between impulse responses of 95-99 percent group and 90-95 percent group (left column); differences between impulse responses of 99-100 percent group and 90-95 percent group (right column). Changes relative to average lagged net income at segments of the liquid asset distribution.

The model estimated is given by $(y_{i,t+h} - y_{i,t-1})/inc_{i,t-1} = \delta_{i,l}^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{l=2,3} \beta_l^h D_{i,t-1}^l \cdot \epsilon_t^{MP} + \sum_{l=1,2,3} \sum_{k=1}^K \gamma_{l,k}^h X_{i,t-k} D_{i,t-1}^l + u_{i,t}^h$ for all i in groups $l = 1, 2, 3$, where $D_{i,t-1}^l$ equals 1 if i is in l in $t - 1$. Groups considered are percentiles 90-95 ($l = 1$), 95-99 ($l = 2$), and 99-100 ($l = 3$). Shown are β_2^h and β_3^h with 95 and 68 percent confidence bands, computed using Driscoll and Kraay (1998) standard errors.

C.5 Net Interest Rate Exposure

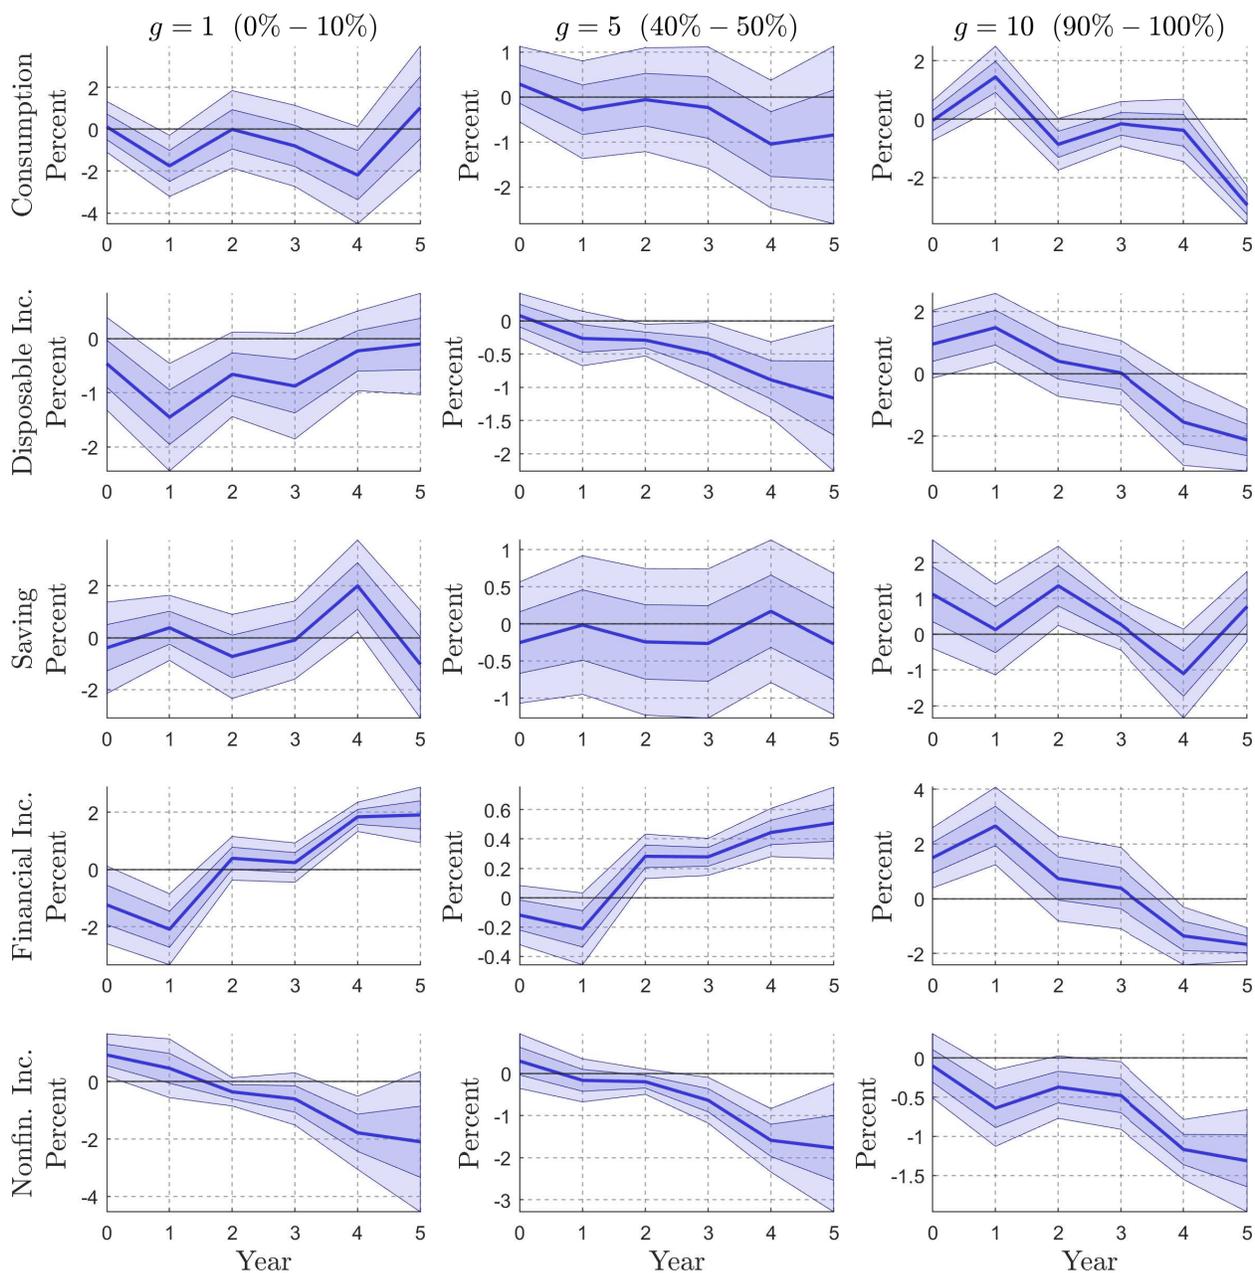


Figure 60: **Impulse Responses of Selected Groups by Net Interest Rate Exposure.**

Notes: Changes relative to lagged net income at segments of the distribution of net interest rate exposure. 95 and 68 percent confidence bands shown, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

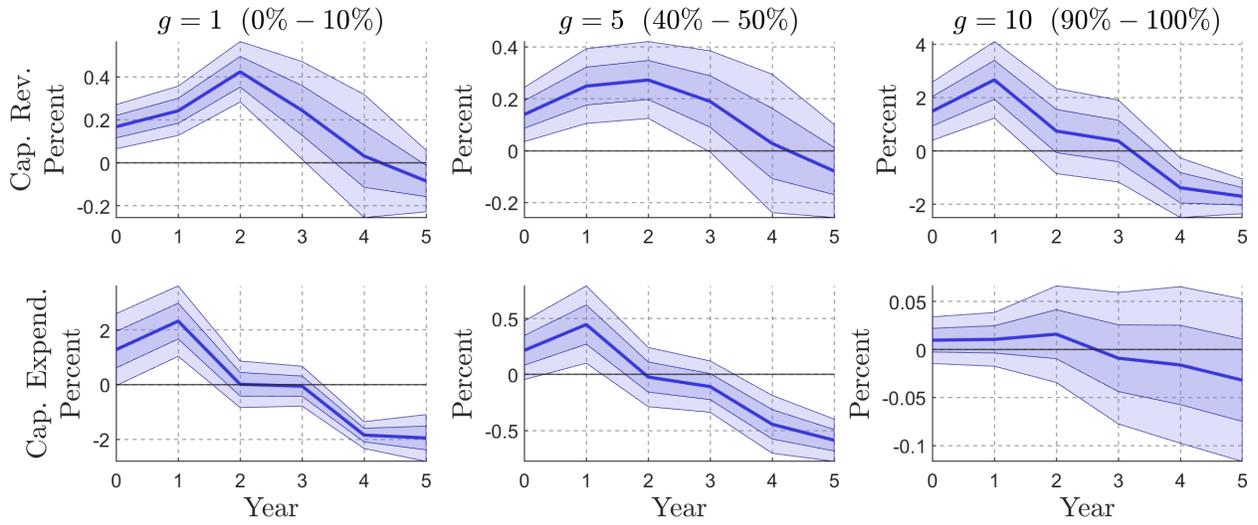


Figure 61: Impulse Responses of Selected Groups by Net Interest Rate Exposure (Continued).
Notes: Changes relative to lagged net income at segments of the distribution of net interest rate exposure. 95 and 68 percent confidence bands shown, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

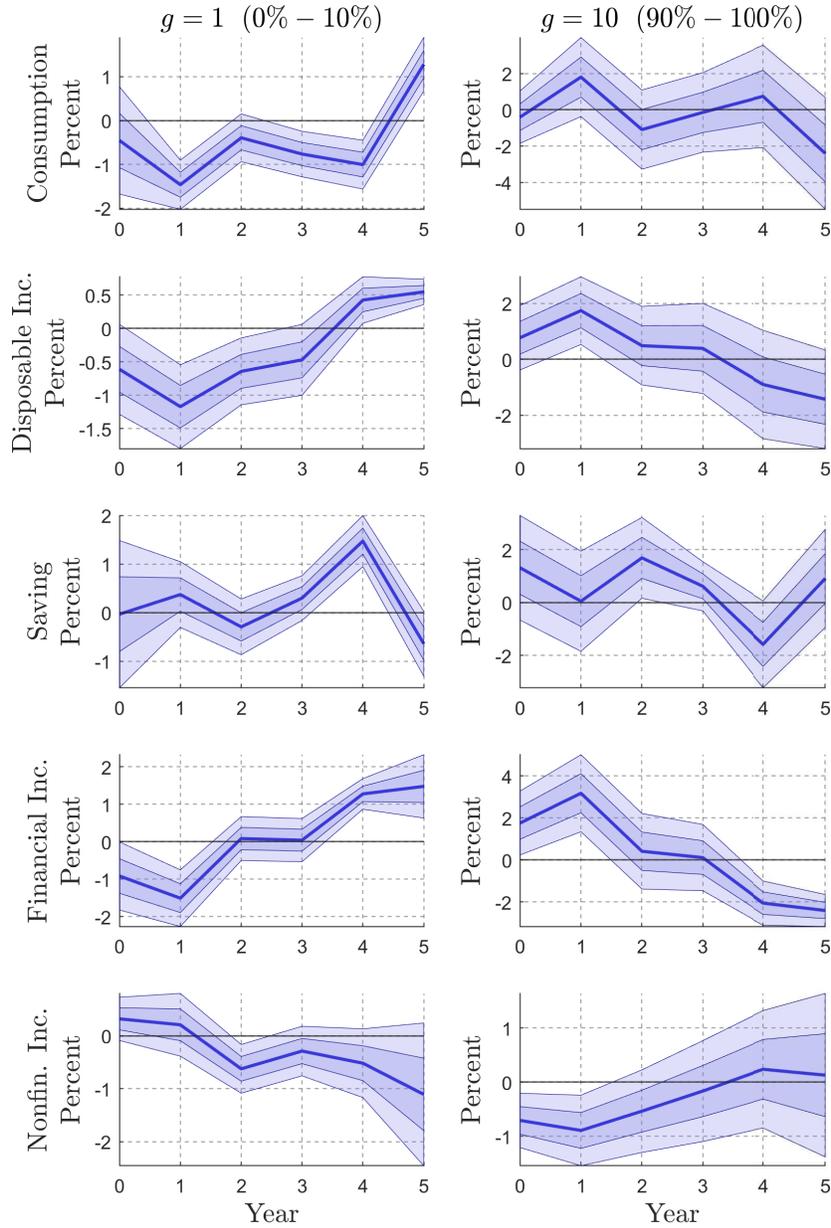


Figure 62: Impulse Responses at Bottom and Top Relative to Group at Median by Net Interest Rate Exposure.

Notes: Differences between impulse responses of 0-10 percent group and 40-50 percent group (left column); differences between impulse responses of 90-100 percent group and 40-50 percent group (right column). Changes relative to lagged net income at segments of the distribution of net interest rate exposure.

The model estimated is given by $(y_{i,t+h} - y_{i,t-1})/inc_{i,t-1} = \delta_{i,g}^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{g=1,10} \beta_g^h D_{i,t-1}^g \cdot \epsilon_t^{MP} + \sum_{g=1,5,10} \sum_{k=1}^K \gamma_{g,k}^h X_{i,t-k} D_{i,t-1}^g + u_{i,t}^h$ for all i in groups $g = 1, 5, 10$, where $D_{i,t-1}^g$ equals 1 if i is in g in $t - 1$. Shown are β_1^h and β_{10}^h with 95 and 68 percent confidence bands, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

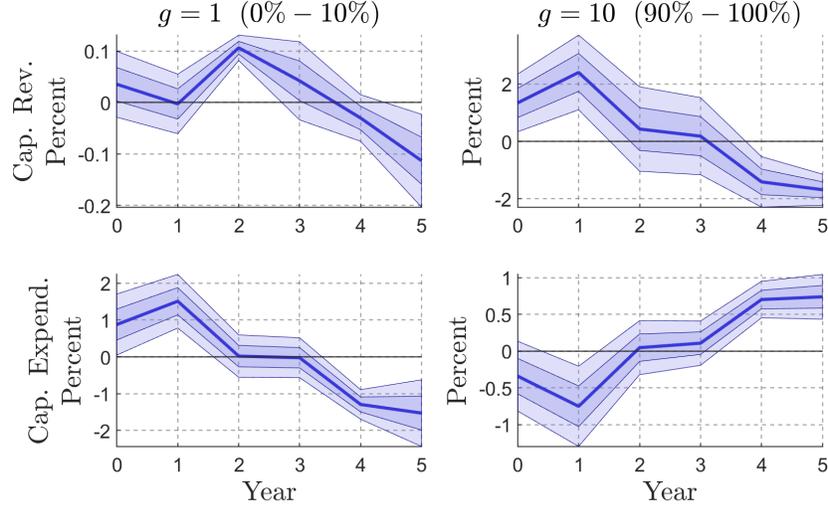


Figure 63: **Impulse Responses at Bottom and Top Relative to Group at Median by Net Interest Rate Exposure (Continued).**

Notes: Differences between impulse responses of 0-10 percent group and 40-50 percent group (left column); differences between impulse responses of 90-100 percent group and 40-50 percent group (right column). Changes relative to lagged net income at segments of the distribution of net interest rate exposure.

The model estimated is given by $(y_{i,t+h} - y_{i,t-1})/inc_{i,t-1} = \delta_{i,g}^h + \beta^h \cdot \epsilon_t^{MP} + \sum_{g=1,10} \beta_g^s D_{i,t-1}^s \cdot \epsilon_t^{MP} + \sum_{g=1,5,10} \sum_{k=1}^K \gamma_{g,k}^h X_{i,t-k} D_{i,t-1}^s + u_{i,t}^h$ for all i in groups $g = 1, 5, 10$, where $D_{i,t-1}^s$ equals 1 if i is in g in $t - 1$. Shown are β_1^h and β_{10}^h with 95 and 68 percent confidence bands, computed using [Driscoll and Kraay \(1998\)](#) standard errors.

D Model

This section lays out the effects of a monetary policy shock on households with different levels of liquid asset holdings in a standard consumption-saving model with a borrowing constraint.

Households are infinitely-lived and have access to a risk-free bond, which they can use to insure against idiosyncratic income risk. Period utility is of CRRA form with elasticity of intertemporal substitution $1/\gamma$. They maximize the expected utility from consumption c , discounted at rate ρ ,

$$\max_{\{c_t\}_{t \geq 0}} \mathbb{E}_0 \int_0^{\infty} e^{-\rho t} \frac{c_t^{1-\gamma}}{1-\gamma} dt \quad (10)$$

Household wealth a_t evolves according to

$$da_t = (r_t a_t + w_t z_t - c_t) dt, \quad (11)$$

where r_t is the (real) interest rate, w_t is the wage, and z_t is the idiosyncratic component of income. z_t evolves according to an Ornstein-Uhlenbeck process given by

$$dz_t = -\nu(z_t - \bar{z})dt + \sigma dW_t \quad (12)$$

ν is the rate of convergence, σ is diffusion, and W_t is a Wiener process. Households face an exoge-

nous borrowing constraint \underline{a} , i.e.

$$a_t \geq \underline{a} \tag{13}$$

Figure 64 depicts the model predictions following an unanticipated 1 percentage point increase in the interest rate for two types of households. High-liquidity households are savers with initial bond holdings at the amount of two times their annual income. Low-liquidity households are borrowers with initial debt of about ten percent of their annual income, which puts them just above the borrowing constraint. The parameterization chosen, $\gamma = 10$, $\rho = 0.041$, $r = 0.04$, $\sigma = 0.1$, and $\nu = 0.03$, is standard with the exception of a relatively low elasticity of intertemporal substitution (high γ) that we select to allow the income effect to be large compared to the substitution effect.

Following the increase in the interest rate, the high-liquidity household experiences a rise in interest income. The change in disposable income triggers a saving response of nearly the same magnitude though, implying that consumption remains almost unaffected. Even for significantly higher values of γ , the saving response is at least an order of magnitude larger than the consumption response. In contrast, the low-liquidity household sees an increase in interest expenses. Since it is close to the borrowing constraint, it has to adjust consumption. The result is that almost all of the disposable income change of the low-liquidity household feeds into consumption through the cash-flow channel.

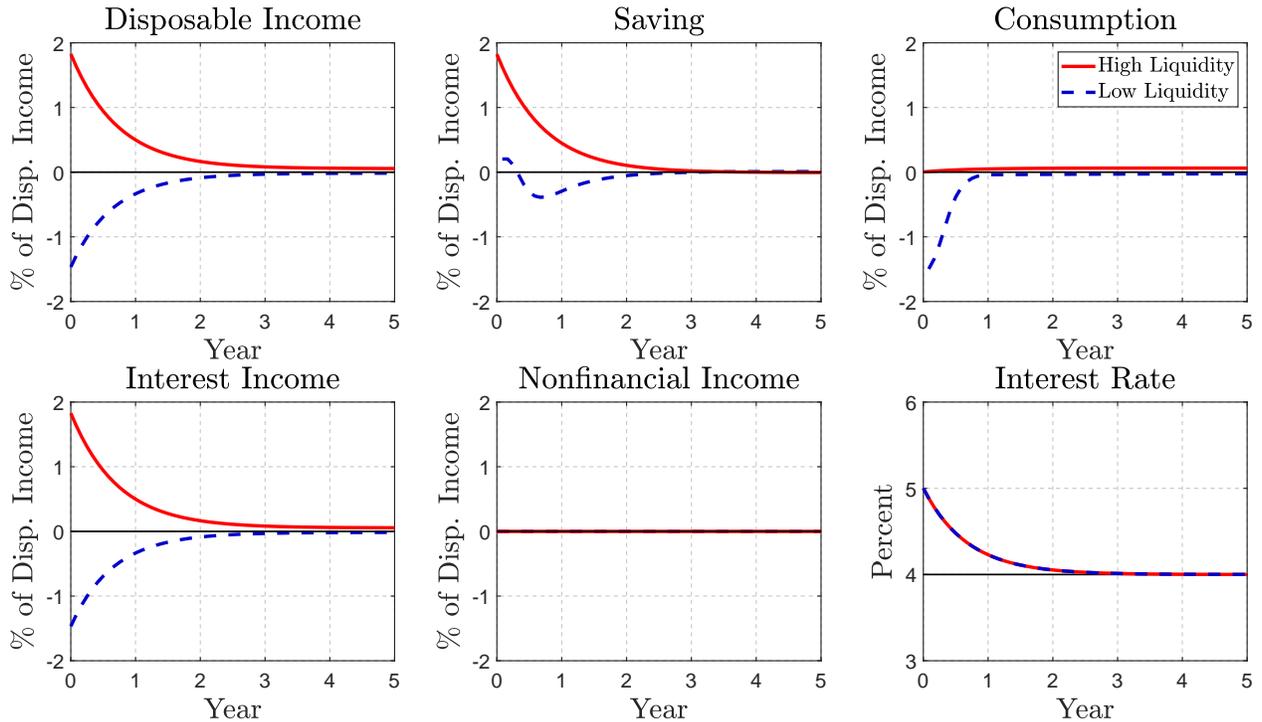


Figure 64: **Impulse Responses in Consumption-Saving Model.**

Notes: Impulse responses to a 1 percentage point increase in the interest rate. Initial wealth is two times annual income (“High Liquidity”) and about -10 percent of annual income (“Low Liquidity”), respectively. The simulation is performed by comparing paths in (i) a benchmark case with no change in the interest rate and (ii) a case with the interest rate change. The parameterization is $\gamma = 10$, $\rho = 0.041$, $r = 0.04$, $\sigma = 0.10$, $\nu = 0.03$.

E Direct and Indirect Effects

E.1 Interpretation and Identification

In this section, we show that the decomposition in Section 5 can be interpreted as a typical omitted variable bias problem in empirical work. To gather intuition, and without loss of generality, we simplify the notation. Denote the local projections in (4) and (7) for consumption as

$$\Delta^h C_{t-1} = \alpha^h + \beta^h \epsilon_t^{MP} + u_t^h \quad (14)$$

$$\Delta^h C_{t-1} = \tilde{\alpha}^h + \tilde{\beta}^h \epsilon_t^{MP} + \gamma^h \Delta^h Y_{t-1} + \bar{u}_t^h, \quad (15)$$

where β^h in (14) gives the response of consumption from $t - 1$ to some $t + h$, denoted $\Delta^h C_{t-1}$, to a monetary policy shock ϵ_t^{MP} at time t , capturing both direct and indirect effects. Regression (15) replicates (7) and differs from (14) by controlling for the change in income $\Delta^h Y_{t-1}$ over the same impulse response horizon. For simplicity, equations (14) and (15) omit all other controls and all household-specific notation. As explained in Section 5, one can interpret $\tilde{\beta}^h$ as the direct effect of monetary policy. Next, consider the auxiliary regression

$$\Delta^h Y_{t-1} = \delta_0^h + \delta^h \epsilon_t^{MP} + \bar{u}_t^h, \quad (16)$$

where δ^h gives the response of income to a monetary policy shock. One can then show that

$$\begin{aligned} \beta^h &= \text{var}(\epsilon_t^{MP})^{-1} \text{cov}(\epsilon_t^{MP}, \Delta^h C_{t-1}) \\ &= \tilde{\beta}^h + \text{var}(\epsilon_t^{MP})^{-1} \text{cov}(\epsilon_t^{MP}, \Delta^h Y_{t-1}) \cdot \gamma^h \\ &= \tilde{\beta}^h + \delta^h \cdot \gamma^h, \end{aligned}$$

such that the total effect β^h can be decomposed into a direct effect $\tilde{\beta}^h$ and an indirect effect $\delta^h \cdot \gamma^h$, which consists of the response of income to the shock δ^h and the partial effect of income on consumption γ^h . Here, $\delta^h \cdot \gamma^h$ can be interpreted as the bias of β^h from the omission of the income controls (see, e.g., Section 2.24 in Hansen, 2019). In fact, this decomposition closely resembles the one in Kaplan, Moll and Violante (2018) (see, e.g., equation (3) in their paper).

To achieve identification of $\tilde{\beta}^h$ using (15), two conditions have to be satisfied. First, income cannot respond to a change in consumption within the same period. For example, that would be the case if a household lowers its consumption and increases its saving, and additionally consumes out of the additional interest income within the same period. However, since interest income and expenses are excluded from the set of income controls, any bias from such a reverse relation is likely to be small within our setup. Second, the structural variation that drives household-level income should not have a contemporaneous direct impact on consumption, unless one controls for such variation in (15). For example, if idiosyncratic labor income shocks are the only variation that drive household income aside from the monetary policy shock, then (15) identifies the true $\tilde{\beta}^h$.

E.2 Instrumental Variable Estimation

To ensure that the second condition is satisfied, we test the robustness of our baseline results in Figure (19) by using an instrumental variable regression. To this end, we use lottery winnings as instruments for household idiosyncratic labor income variation based on the Norwegian lottery data from [Fagereng, Holm and Natvik \(2018\)](#). These data contain all individuals who have won a prize greater than USD 1,100 (USD 11,000 after 2006) in a game administered by Norsk Tipping, the state-run gambling monopoly in Norway. All games administered by Norsk Tipping include a significant element of luck such as bingo, scratch cards, sports betting, horse racing, and lotto. There are approximately 30,000 unique lottery winners within our sample.⁴²

Continuing with the simplified notation above that omits our standard controls and any household-specific notation, we estimate the first-stage regression

$$\Delta^h Y_{t-1} = \tilde{\delta}_0^h + \tilde{\delta}_1^h \epsilon_t^{MP} + \tilde{\delta}_2^h \text{lottery}_{t+h} + \tilde{u}_t^h, \quad (17)$$

where the dependent variable is household nonfinancial income and lottery_{t+h} is the lottery winning in period $t + h$.⁴³ Since lottery winnings enter into the income definition in (17), we obtain estimated coefficients $\tilde{\delta}_2^h$ that are close to one and statistically significant with p-values below 0.01 for all impulse response horizons. In a second stage, we use the predicted values $\Delta^h \hat{Y}_{t-1}$ as income controls in (15). Intuitively, the lottery winnings offset the impact of a monetary policy shock in (15) that is running through income and allow to hold income constant to estimate $\tilde{\beta}^h$. The results based on this IV-setup are shown in Figure (20).

E.3 Accounting for Indirect Wealth Effects

To account for indirect wealth effects, we add measures of changes in households' risky asset wealth into the specification in (7). In fact, the estimated capital gains as described in Section B.1 directly capture such wealth changes. We therefore extend (7) as

$$\frac{c_{i,t+h} - c_{i,t-1}}{\bar{c}_{t-1}} = \delta_i^h + \beta^h \epsilon_t^{MP} + \sum_{m=1}^3 \mu_m^h X_{i,t-m} + \sum_{k=1}^K \sum_{m=0}^h \gamma_m^{h,k} \tilde{y}_{i,t+m}^k + \sum_{m=0}^h b_m^h \frac{CG_{i,m}}{\bar{c}_{t-1}}, \quad (18)$$

where the new term $\sum_{m=0}^h b_m^h \frac{CG_{i,m}}{\bar{c}_{t-1}}$ captures capital gains for household i between period $t - 1$ and $t + h$, scaled by $t - 1$ per capita consumption. The results based on (18), compared against the consumption response without income and capital gain controls, are shown in Figure 65. Indirect wealth effects turn out to be quantitatively small, such that the results in Figures 19 and 65 are similar.

⁴²We exclude winners of multiple prizes and trim the size of the lottery winnings to be below 50,000 USD to account for the fact that monetary policy shocks are unlikely to result in larger income changes.

⁴³As in regression (7), we scale the dependent variable, the income controls, and the lottery winnings by average consumption \bar{c}_{t-1} across all households at time $t - 1$. We further control for all income changes separately from $t - 1$ to any $t + k$ where $k = 0, \dots, h$.

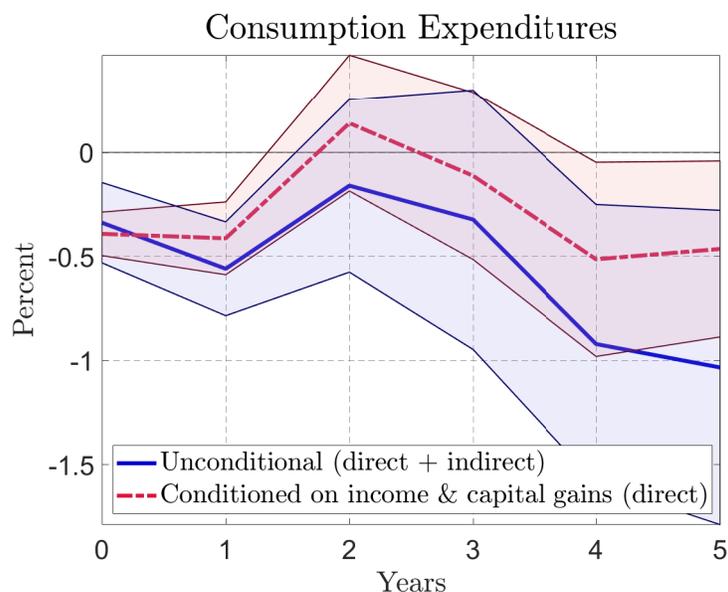


Figure 65: **Direct and Indirect Effects of Monetary Policy.**

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 (18). The blue line shows the estimated impulse responses without controlling for income, the red dashed line shows the responses with income and capital gains controls. 68 percent confidence bands shown, using Driscoll and Kraay (1998) standard errors.

E.4 Direct and Indirect Effects Across the Liquid Asset Distribution

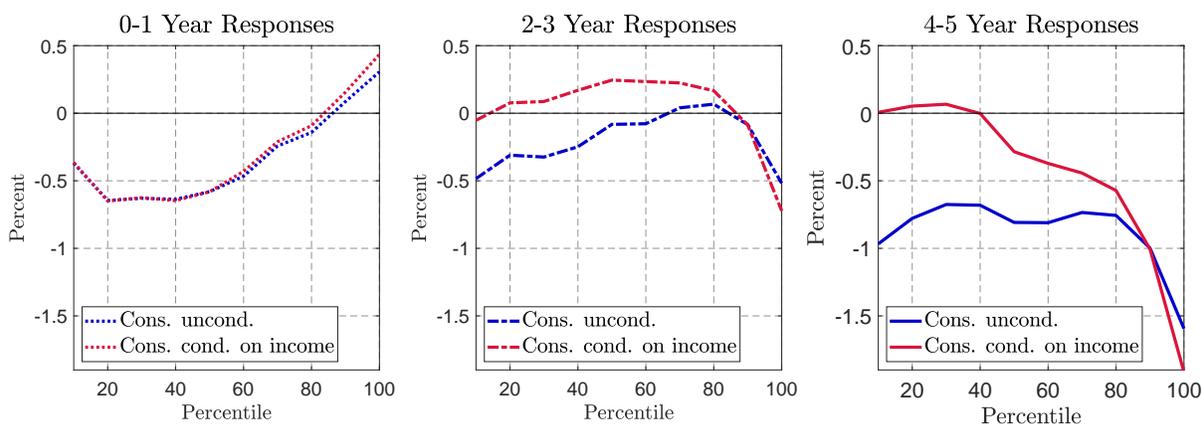


Figure 66: **Direct and Indirect Effects of Monetary Policy.**

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 (7) for decile groups along the liquid asset distribution. The blue lines show the estimated impulse responses without controlling for income, the red lines show the responses with income controls.

E.5 MPC Estimates

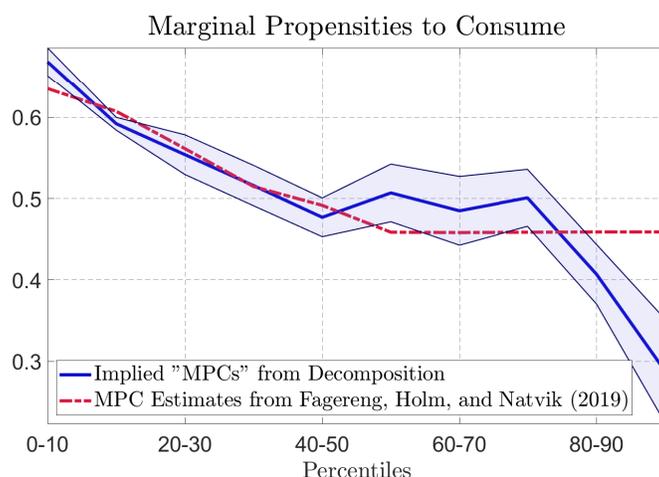


Figure 67: **Implied and Estimated MPCs.**

Notes: The blue line shows the estimated coefficient associated with the contemporaneous nonfinancial income control for the time zero impulse response horizon $\hat{\gamma}_0^{0,k}$ based on regression (7) across the liquid asset distribution. The red dashed line shows interpolated MPC estimates from Fagereng, Holm and Natvik (2018) for different levels of liquidity. There are fewer observations in the sample by Fagereng, Holm and Natvik (2018) for high levels of liquidity, resulting in a relatively flat line above around the 60th percentile.

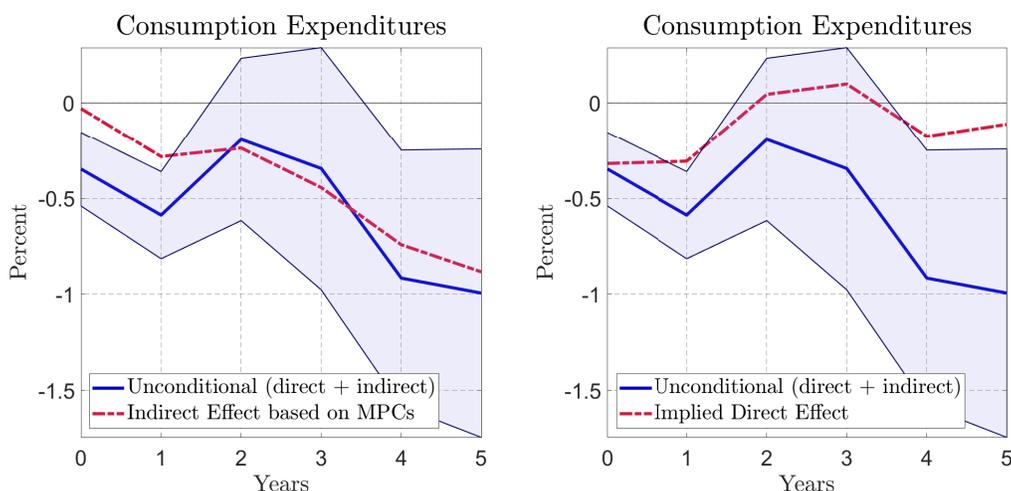


Figure 68: **Direct and Indirect Effects of Monetary Policy – Decomposition based on MPCs.**

Notes: The blue lines show the estimated impulse response of consumption expenditures based on the local projection approach in (4) that is shown in Figure 9. The red dashed line in the left graph gives an estimate of the indirect effect. It is computed as the product of the dynamic MPC estimates from Fagereng, Holm and Natvik (2018) (see Section 3.2. therein) and the impulse response of disposable income shown in Figure 9. The dynamic MPCs are $MPC_{t,t} = 0.5057$, $MPC_{t+1,t} = 0.1759$, $MPC_{t+2,t} = 0.1035$, $MPC_{t+3,t} = 0.0444$, $MPC_{t+4,t} = 0.0337$, and $MPC_{t+5,t} = 0.0255$, where $MPC_{t+k,t}$ denotes the marginal propensity to consume at time $t+k$ to a transitory income shock at time t . The red dashed line in the right graph is the difference between the two lines in the left graph.