The Cash-Flow Channel of Monetary Policy - Evidence from Billions of Transactions*

SeHyoun Ahn
Norges Bank

Sigurd Mølster Galaasen
Norges Bank

Mathis Mæhlum
Norges Bank

This version: July 8, 2024

Abstract

We present novel findings on the impact of monetary policy on consumer spending behavior using a recently assembled high-frequency household expenditure panel. Leveraging comprehensive weekly electronic transaction-level data for all individuals in Norway over 13 years, our study sheds light on the high-frequency consumption response to monetary policy through changes in fast-moving debt interest expenses. We employ several identification strategies, including household–specific interest rate shocks arising from a natural experiment as well as estimated monetary policy shocks. Our results show a substantial short-run marginal propensity to consume out of interest payments. We find that 100 USD higher interest expenses lead to around 30 USD lower consumption by the end of the first year, and that the effect materializes gradually throughout the year. This suggests the presence of a strong cash-flow channel of monetary policy operating through fast-moving interest expenses.

Keywords: Monetary policy transmission, Household Consumption, Debt, Interest rates, Marginal propensity to consume, High frequency

JEL Classification: D12, D31, E21, E43, E52, G51

*Ahn: Norges Bank, sehyoun.ahn@norges-bank.no. Galaasen: Norges Bank, sigurd-molster.galaasen@norges-bank.no. Mæhlum: Norges Bank, mathis.mehlum@norges-bank.no. We are grateful to Karsten Gerdrup and Kjersti Næss Torstensen for their invaluable contributions to collecting the electronic payments data and for their insightful discussions. We also thank Lars E.O. Svensson, Silvia Miranda-Agrippino, Federica Romei, Martin Blomhoff Holm, Saskia Ter Ellen and seminar participants at SNDE Annual Conference 2024, IAAE Annual Conference 2024, 10th edition of Conference on New Developments in Business Cycle Research (Danmarks Nationalbank), Sveriges Riksbank, Statistics Norway, Norwegian School of Economics and Norges Bank for invaluable comments and suggestions. This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank.

1 Introduction

What is the role of household debt in the transmission of monetary policy? In recent decades, household debt has increased faster than income in many developed economies, and in several countries the average debt-to-income ratio is now at an historically high level. While a growing literature (Mian et al., 2013; Jordà et al., 2016) has shown that household leverage can amplify macroeconomic fluctuations, whether increasing household indebtedness also influence Central Banks' ability to stabilize aggregate demand remains an open question. Since monetary policy transmits rapidly to mortgage rates for holders of adjustable rate mortgages, rising debt-to-income ratios means that a greater share of household cash flows become directly exposed to policy changes, potentially affecting the speed and size of household demand responses. In this paper, we offer empirical evidence on how debt influences monetary policy transmission by examining how the responsiveness of household consumption to interest rates varies with interest exposure.

Traditionally, the standard macroeconomic view has been that cash-flow fluctuations have only a minimal impact on consumption. Rational agents with frictionless access to capital markets adjust consumption smoothly in response to monetary policy-induced income changes. In the workhorse New Keynesian models (e.g. Clarida et al. (1999) and Christiano et al. (2005)), monetary policy mainly transmits to consumption through the intertemporal substitution channel, while income effects are small.[1] In contrast, recent theoretical (Kaplan et al., 2018) and empirical (Di Maggio et al. (2017); Flodén et al. (2020); Cloyne et al. (2020)) advances suggest that households' balance sheet positions are pivotal in transmitting monetary policy shocks to households consumption. Nevertheless, the empirical understanding of the role of debt in transmitting monetary policy remains limited, primarily due to demanding data requirements. The ideal setting, which rarely exists, requires high-frequency household-level consumption expenditures linked to detailed balance sheet positions over time.

[1] This is demonstrated by (Kaplan et al., 2018), among others.

The Norwegian economy, with its comprehensive administrative records, provides a unique opportunity to address this data challenge. In this paper we combine household balance sheet and income information from high quality tax records with a newly assembled high-frequency electronic expenditure database. Our dataset contains detailed information on wealth, debt, income, and demographics, merged with weekly debit card purchases and online bank wire transactions for nearly every individual in Norway between 2006 and 2018. As Norway is virtually a cashless economy, the electronic expenditures provide a comprehensive measure of household-level consumption. Its high-frequency nature also enables measurement of the short-term responses to monetary policy and makes it possible to overcome a potential temporal aggregation bias by aligning the timing of the shock with the outcome variable, as discussed in Buda et al. (2023) and Jacobson et al. (2022).

With the rich micro level data at hand, the paper proceeds with investigating the role of debt for the transmission of monetary policy to household consumption. We do so by estimating how household level variation in interest payments, induced by interest rate changes, transmits to consumption. In order to address the inherent endogeneity associated with interest rate changes, we approach the research question using several identification strategies. They produce similar results.

Following a large empirical macro literature studying the impact of monetary policy (see e.g. Ramey (2016)), the first two approaches employ local projection methods à la Jordà (2005) to estimate the dynamic impact of a change in the policy rate on household-level consumption at a monthly frequency. We interact movements in the policy rate with household interest exposure (debt minus deposits), creating a proxy for changes in interest payments. To address endogeneity issues arising from monetary policy reacting to changes in macroeconomic conditions, we use two separate local projection specifications. In the first specification, we include time fixed effects interacted with granular group level characteristics. These fixed effects are intended to account for the confounding macroeconomic environment so that we can separately identify the effect of varying interest exposure on consumption. In the second approach, we construct a new set of monetary policy instruments for the Norwegian economy and use them in an instrumental variable setup as in Stock and Watson (2018). We partition households into bins based on their position in the interest exposure-to-income distribution and estimate local projection instrumental variable (LP-IV) regressions for each group and time horizon separately. We instrument the interest rate movements with high-frequency monetary policy surprises in the spirit of e.g. Kuttner (2001) and Gürkaynak et al. (2004), and we follow Miranda-Agrippino and Ricco (2021) in removing a potential information component.

After estimating how the transmission of monetary policy to consumption varies with interest exposure and debt, we proceed by interpreting the results as marginal propensities to consume out of changes in interest payments. We do so by assuming that the interaction of pre-determined household debt and a one percentage point change in interest rates effectively measures the change in interest rate payments households experience. The marginal propensity to consume is then inferred from the size of the consumption change relative to the proxy for interest payment changes.

In the third approach, in contrast to the first two identification methods, we measure changes in interest payments resulting from household-specific shocks to mortgage rates arising from a natural experiment unrelated to macroeconomic conditions. The experimental setting offers two main benefits compared to the earlier methods, alleviating potential measurement and endogeneity concerns. First, rather than relying on a proxies, we now observe both interest payments and borrowing rates directly in the data. Second, the identifying variation in interest expenses arises from changes in the cross-sectional interest rate dispersion, rather than through dispersion in debt combined with aggregate interest rate movements. The natural experiment involves a mortgage bank serving exclusively public sector workers, which prior to March 2014 offered highly subsidized mortgage rates relative to those offered in the private sector mortgage market. For the 2014 and 2015 National Budget, the government decided to cut this subsidy heavily, inducing a sharp rise in mortgage rates in both March 2014 and subsequently in March 2015 for clients of the public sector bank only. Over this period, the subsidized rate was still lower than the rates offered in the conventional mortgage market, so there was no incentive for the customers of the public sector bank to switch. Using detailed information on household bank connections, we then compare the evolution of consumption and interest payments for households with mortgages in the public sector bank with customers of other banks. This allows us the estimate the MPC out of interest expenses.

From the local projections of consumption changes on interest rate changes, we observe that the impact of higher interest rates is more pronounced for households with higher levels of debt. The impact is both rapid and quantitatively important: from the fixed-effects specification, we find that a doubling of the interest exposure-to-income ratio is associated with a roughly 0.2 percentage points larger drop in consumption relative to income. This effect appears gradually throughout the first six months following a monetary policy change. Splitting the households into bins and instrumenting for monetary policy, we find that this effect is present along most of the interest exposure distribution, suggesting that it is not solely driven by households with exceptionally high levels of debt.

However, among households with low and negative levels of interest exposure – due to a combination of low debt and a substantial bank deposits – the estimated consumption responses do not vary with exposure. Thus, there is a distinct non-linear pattern at the lower end. This indicates that households react to changes in borrowing rates, but not deposit rates.

A back-of-the-envelope calculation based on the above results implies that when interest payments increase by 100 USD, consumption falls on average by around 20 USD after 6 months and 30 USD after 1 year. The consumption response of households who hold no debt is close to zero, while households at the 90th percentile of debt-to-income experience a consumption fall corresponding to around 2 percent of income after 1 year. This implied MPC is confirmed when we employ the natural experiment. Here we find that, compared to households that are similar among other dimensions, customers in the public sector bank saw a large relative increase in interest payments of 912 USD in 2014. These changes are entirely driven by the relative changes in mortgage rates induced by the 2014 policy change. In response to these movements in interest expenses, we find that the public sector bank customers reduced their consumption by 278 USD in 2014 relative to comparable households, implying an average MPC of 30 percent.

Our paper contributes to several strands of the literature. First, by all accounts our results confirms the presence of sizable average MPCs also out of interest payment shocks, in line with recent evidence from other types of income shocks (Parker et al., 2013; Jappelli and Pistaferri, 2014; Parker, 2017; Carroll et al., 2017; Ganong et al., 2020; Lewis et al., 2019; Aguiar et al., 2020; Andersen et al., 2021; Gerard and Naritomi, 2021; Fagereng et al., 2021; Gelman, 2022; Boehm et al., 2023; Hamilton et al., 2023). Our results also indicate a potential asymmetry across income and expenses, as we find muted consumption responses for people with little debt, but sizable deposits.

Second, we contribute to the literature estimating the transmission of monetary policy to household consumption at the micro level. In particular, we provide estimates of the interest rate exposure channel described theoretically by Auclert (2019). Like Flodén et al. (2020), Cloyne et al. (2020), Holm et al. (2021), Di Maggio et al. (2017) and La Cava et al. (2016), we find an important role for mortgage debt in the transmission of unexpected monetary policy shocks to consumption. Unlike prior studies that utilize either survey data (Cloyne et al. (2020) and Wong et al. (2019)), impute consumption expenditures from annual tax returns data (Holm et al. (2021) and Flodén et al. (2020)), or use a subset of expenditures (Di Maggio et al. (2017)), we construct a consumption variable that is directly measured, captures most of a household's spending and is observed at a high frequency that aligns with the timing of monetary policy shocks. As such, we also contribute to the literature that estimates the high-frequency responses of aggregate variables to monetary policy shocks. Like Buda et al. (2023), we find that changes in monetary policy can transmit to household consumption more quickly than many previous studies utilizing quarterly or annual variables have found.

Third, our study can potentially shed light on the role played by housing markets, mortgage markets and mortgage debt in shaping the transmission of monetary policy to the real economy. International Monetary Fund (2024) finds that transmission is stronger in countries where, like in Norway, fixed rate mortgages are not common, where home buyers are more leveraged and where household debt is high. Similar results are reported by Pica (2023). Alpanda et al. (2021) find that the effect of monetary policy on GDP is stronger in periods when the household debt-to-GDP ratio is above its long-run trend, but only in countries where adjustable rate mortgages are the norm. Our results quantifies this channel of transmission at the micro level.

The paper proceeds as follows. In Section 2 we present the our linked expenditure data and administrative records. In Section 3 we estimate how the response to common interest rate changes vary with household debt, and impute the marginal propensity to consume out of interest payments. In Section 4 we utilize the natural experiment providing us with household specific interest changes, allowing us to obtain a direct estimate of the MPC. We summarize our findings in Section 5.

2 Data

We combine macroeconomic data with individual level register data from Statistics Norway and electronic expenditure data provided by Nets Branch Norway. This section provides an overview of each of the data sources and explain the criteria used to construct our analysis sample.

2.1 Expenditure Data From Electronic Transactions

This section outlines in detail our electronic payments database, which includes most card and bank wire transactions for all Norwegian residents. We start by explaining how the database was collected and organized in 2.1.1. Then, in 2.1.2, we discuss the payment coverage and show that it contains about 80 percent of all electronic payments made by Norwegian households over the time period 2006-2018. In 2.2.1 we compare our expenditure measure with household consumption in the National Accounts.

2.1.1 Data Collection and Structure

The payments data are provided by the Norwegian retail clearing institution, Nets Branch Norway (henceforth abbreviated by Nets). The data covers two types of payments: (i) debit card payments made via BankAxept,[2] and (ii) bank wire transfers cleared via the Norwegian Interbank Clearing System (NICS).[3] The data includes all Norwegian individuals that have made a debit card transaction using BankAxept or bank wire transfer via the NICS clearing system in the period 2006 to 2018.

The transactions are aggregated by the data provider to the level of person, week, postal code and consumer category. This aggregation, explained further below, is based on underlying metadata stored with each transaction by Nets. Table **??** presents all variables in the raw data provided by Nets to Norges Bank. In addition to the number and amounts of transactions per person per week, consumer category and location, each observation includes a set of individual demographic characteristics: birth year and gender. For debit card transactions, the data also includes a variable recording cash withdrawals made during the purchase (cashback). However, ATM cash withdrawals are not included in the data.

The consumption categorisation is based on 24 COICOP groups,[4] and contains all 12 top level codes such as "Food and beverages", "Restaurants and hotels" and "Clothing and footwear" and "Housing, water, electricity, gas". Some top-level categories are further divided into second-level COICOP groups. Additionally there are two categories (13) and (14) which apply only to bank wire transfer, bringing the the total to 26 categories. Table 1 lists all categories.

The aggregation by geography and consumer categories is based on seller type and location information. The consumer categorisation is based on the United Nations' 1999

[2]BankAxept serves as a national payment processing system in Norway. In the time period 2006-2018 debit card transactions at physical domestic stores typically utilized BankAxept, while debit card payments made abroad, online, or via mobile platforms were processed through VISA or Mastercard.

[3]NICS is the interbank clearing system for the Norwegian Krone (NOK) and is used by all banks operating in Norway and that take part in the Norwegian banking community's infrastructure for payments.

[4]COICOP is an abbreviation for Classification of Individual Consumption According to Purpose, developed by the United Nations Statistics Division.

Table 1: Consumption categories

01	Food and non-alcoholic beverages		
02	Alcoholic beverages, tobacco and narcotics		
03	Clothing and footwear		
04	Housing, water, electricity, gas and other fuels		
05	Furnishings, household equipment and routine household maintenance		
06	Health		
07	Transport		
		071	Purchase of vehicles
		072	Operation of personal transport equipment
		073	Transport services
08	Communications		
09	Recreation and culture		
		091	Audio-visual, photographic and information processing equipment
		092	Major durables for outdoor recreation
		093	Other recreational items and equipment, gardens and pets
		094	Recreational and cultural services
		095	Newspapers, books and stationery
10	Education		
11	Restaurants and Hotels		
		111	Restaurants
		112	Hotels
12	Miscellaneous goods and services		
		121	Personal care
		123	Personal effects
		124	Social protection
		125	Insurance
		126	Financial services
		127	Other services
13	Payments to banks		
14	Payments to public institutions		

Notes: Aggregation level for consumption category. The category numbers corresponds to the 1999 COICOP version. Category 13 and 14 are not part of the COICOP classification, and apply only to bank wire transfers made via NICS.

COICOP classification,[5] As COICOP codes are not part of the metadata that the data provider stored with each individual transaction, we provided Nets with a cross-walk between the metadata and COICOPs to facilitate the aggregation.

For debit card transactions, the metadata contains the date of the transaction, the location (zip code) of the the card terminal, and the business name, NACE Rev. 2 and Merchant Category Code (MCC) of the store. Aggregation by location is based on the four-digit zip code registered on the card terminal,[6] while aggregation by COICOP is based on the MCC code registered with on the card terminal.[7]. Aggregation by COICOP codes is obtained through a cross-walk between MCC and COICOP. Out of 162,000 card terminals, 42,600 were classified with MCC "Miscellaneous and Speciality Retail Stores" (MCC=5999) or had missing MCC. These terminals were classified using the metadata registered on the terminal, typically the name or NACE Rev 2. code (five-digit Norwegian SN07) of the store.[8]

For bank wire transfers, the transaction metadata includes the transaction date, and the creditor's NACE Rev. 2 and location (zip code). Aggregation by location is based on the four-digit zip code, while aggregation by the consumption categories listed in Table 1 is based on a mapping between the NACE Rev 2. code (five-digit Norwegian SN07) and COICOP.[9] Two types of wire transfers, categorized as non-COICOP categories 13 and 14, include transfers between persons and banks (13) and between persons and government institutions (14).

[5] The COICOP classification was revised in 2018. The top level structure (first and second level), the level at which our electronic payments are aggregated, remained mostly unchanged. The main update was the introduction of more granular categories (at the fourth level).

[6] There are about 5000 such codes in Norway. To ensure that the combination of zip-COICOP does not identify expenditures classified as being sensitive (e.g. precise information on type of firm related to health expenditures), location information has been removed (replaced with -1) for certain combinations of COICOP-zip code.

[7] MCC is a four-digit number listed in ISO 18245 for financial services and is used to classify a business according to its main activity.

[8] The MCC-COICOP cross-walk and the classification of undefined (5999) or missing MCC stores by COICOP is available upon request.

[9] The cross-walk is available upon request

2.1.2 Coverage and Cleaning

Debit Card Transactions (BankAxept) In Figure 1, panel (a) and (b), we plot the annual time series of total number of BankAxept transactions and value of payments in the debit card database along with official payment statistics (Norges Bank, 2023b) on BankAxept debit card payments and cashback. The figure shows that from 2012 and onwards our data covers the universe of household debit card transactions. Prior to 2012 the coverage shrinks somewhat. The reason is that Nets Branch Norway did not retain information on bank account owners for accounts that were closed prior to 2012. Consequently, transactions using debit card linked to such accounts to not show up in our database. The likely major source of bank account closing is people who either die or migrate abroad during the time period 2006-2011.

Bank Wire Transferes In Figure 1, panel (c) and (d), we plot annual time series for total number and value of bank wire transfers. From panel (c) we see that our data tracks the official statistics closely, but with a gap. On average our data contains 77 percent of all transfers. The reason we do not have full coverage is that not all bank wire transfers are processed by NICS, in particular wire transfers between owners of accounts in the same bank are sometimes processed by the banks themselves.[10] The jump in transfers occurring in 2018 comes from the introduction of peer-to-peer mobile transfers. These are mobile phone transfers of payments between households. To avoid double counting of consumption expenditures we remove transfers between individuals from our data.[11]

In Panel (d) we see that the coverage of total value of transfers is somewhat lower than for the number of transfers. On average we cover around 60 percent. It is important to emphasize that the value of transfers is highly sensitive to very large single transactions.[12] As these are unlikely to reflect consumption expenditures, we clean our data by removing single transactions with a value above 12,500 USD (100,000 NOK in 2015). In addition,

[10] The official statistics picks up these transfers because they receive data on account-to-account transfers directly from the banks.

[11] This is done by removing all wire transfers with missing information on both COICOP and zip.

[12] In fact, in both 2009 and 2011 there is a single transfers that account for more than 5 percent of total the total transfer value reported in the official statistics.

Figure 1: Transactions and Amounts.

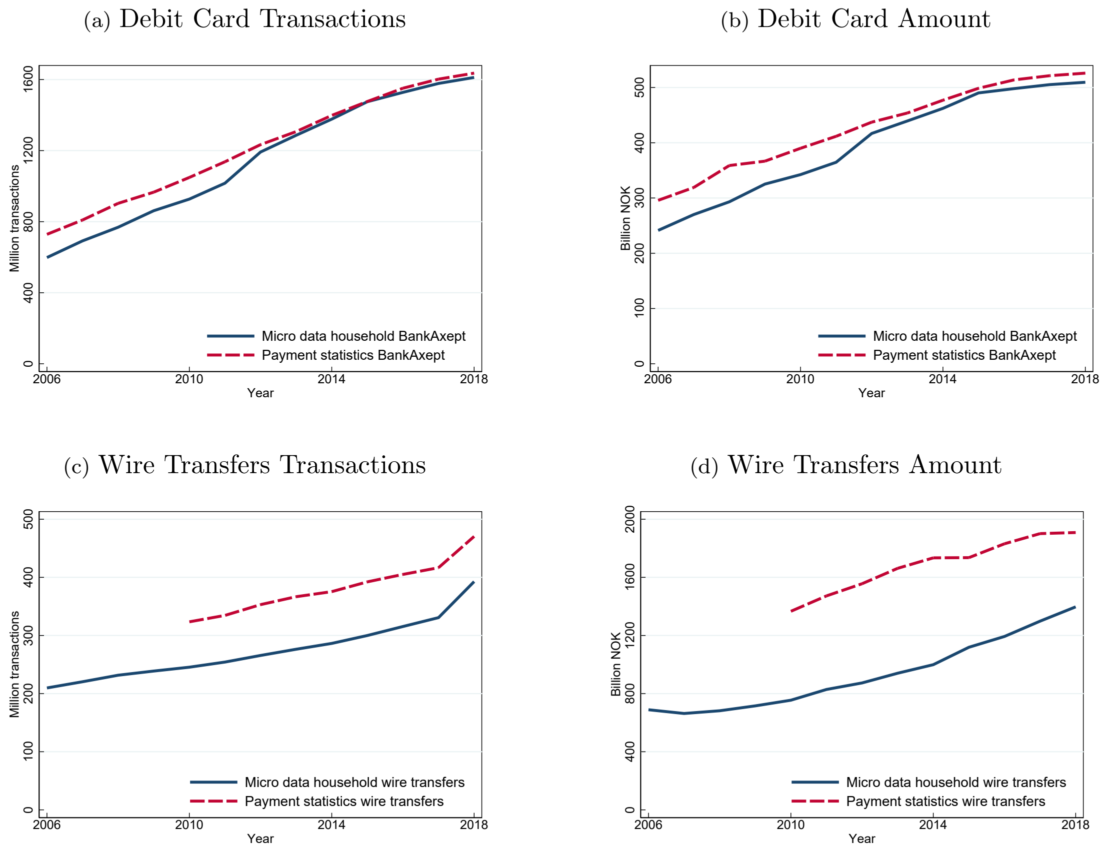

Notes. This figure shows aggregate number of transactions and the total transaction amount from two sources. The solid line represents the aggregated volumes over all households in our data base. The dashed line represents the official numbers reported in Norges Bank (2023b). Panel (a) and (b) show debit card transactions processed using BankAxept. Panel (c) and (d) show bank wire transfers. The split between firm and household wire transfers is only available from 2010 in the official statistics.

all transfers made by households to banks, which we observe in consumption category 13 (see Table 1) are removed from the wire transfer measure. However, as explained below, we retain an imputed measure of payments of credit card bills.

Credit Card Imputation While our card transaction data does not include credit card transactions directly, our wire transfer data encompasses payments made directly to banks. The main source of such payments are related to debt service, including mortgages and credit card bills. In Section A we explain the algorithm we use in order to separate between payments related to credit card bills and those related to regular debt service. The imputed credit card bill payments are retained in our expenditure measure, whereas regular debt service is removed. Figure 2 shows the imputed credit card payments and compare it to official statistics on credit card payments in Norway.

Figure 2: Credit Card Coverage

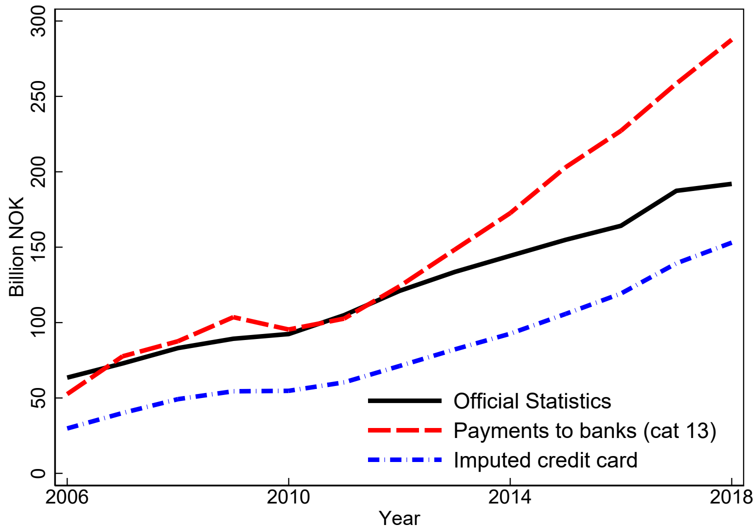

Notes. This figure plots annual time series of official credit card payments (solid line) and total payments in consumption category 13 (dashed line). The dashed-dotted line represents our imputed credit card measure, obtained with the algorithm outline in Section A. Source: Norges Bank (2023b).

Total Card Payments Coverage Lastly, we evaluate the proportion of total card transactions covered by our debit card and imputed credit card payments. We start by showing that debit card transactions processed through BankAxept covers the majority of card payments in Norway. Typically, all debit card payments in domestic physical stores are processed through BankAxept, whereas debit payments abroad, online and mobile payments are processed through VISA or Mastercard. Using aggregate statistics on Norwegian households' card usage, published annually by Norges Bank (Norges Bank, 2023b), Figure 3 shows that BankAxept debit card payments accounts for more than 80 percent of card payments in Norwegian in-store terminals. By adding online payments, which is not covered by BankAxept, the share drops with only a few percentage points, illustrating that online payments was not extensively used during our sample period.

Figure 3: Total Card Coverage.

(a) Debit Card (b) Debit Card and Imputed Credit Card

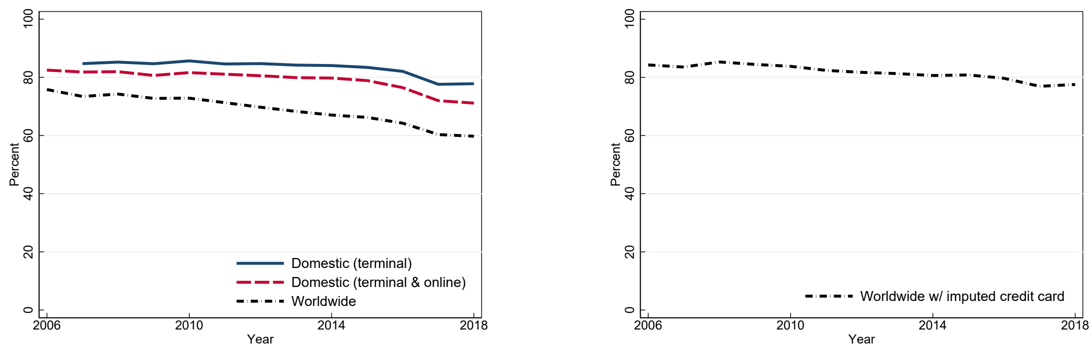

Notes. Panel (a) plots the share of total card payments (number of transactions) in (i) domestic card terminals, (ii) domestic card terminals and online payments and (iii) domestic and foreign terminals and online payments, that is covered by our debit card measure. Panel (b) plots the share of total card payments worldwide covered by our debit card and imputed credit card measure.

Finally, if we consider all card payments made by Norwegian regardless of point of sale, the BankAxept share is on average around 70 percent, with a negative trend over the sample period. This negative trend reflects the increased usage of credit cards. However, when we add our imputed credit card measure, panel (b) shows a stable coverage rate of about 80 percent throughout the sample period.

2.2 Cash payments

While our measure of consumption expenditures does not include payments made by cash, Norway is among the countries with the lowest cash share of payments in the world. In surveys, 3 – 5% of respondents have reported that they used cash for their most recent payment (Norges Bank, 2023a). The cash share of the monetary aggregate M1 was around 6% at the start of our sample and fell to around 2% at the end of our sample (Norges Bank, 2023b).

2.2.1 Comparison with National Accounts Household Consumption

Figure 4 shows the growth rates of our aggregated electronic expenditures and household domestic consumption in the National accounts. Overall, the two series track each other well, and the raw correlation is 0.83.

Figure 4: Comparison with National Accounts

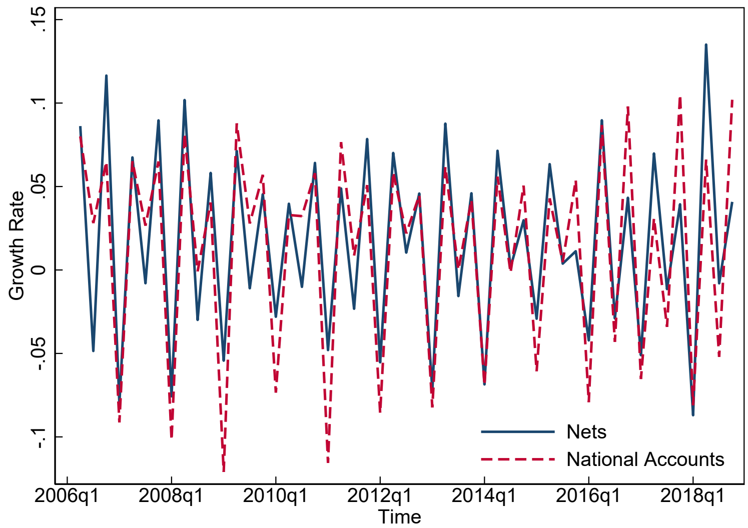

Notes. This figure plots the quarterly growth rates of nominal household domestic consumption from the National Accounts (dashed line) and aggregate household electronic expenditures (solid line). For the expenditure measure we remove single wire transfers above 12,500 USD, person-to-person transfers, and all transfers to banks, except imputed payments of credit card bills.

2.3 Administrative Data

We combine our electronic expenditure data with administrative records containing detailed information on income, assets, demographic characteristics and labor market status, for the universe of Norwegian residents over the period 1993-2018. The income and asset data is collected for tax purpose by the Norwegian Tax Authority, and for the most part comes from pre-filled tax declarations by 3rd party reporting. In particular, the data contains information on income (labor income and capital income and expenses), total debt, financial assets (deposits, securities) and real wealth (housing, cars, business wealth). Income and expense measures represents the flow throughout the year, while assets represents the stock at the end of the year. The demographic information contains among other things socioeconomic characteristics such as age, education and family relations, from which we construct our household measure. The labor market data contains workplace information such as industry and occupation, and comes from employee-employer registers. This data has been extensively used for research, and documented in detail in other studies (see for example, Fagereng et al. (2020) and Holm et al. (2021)). One important aspect of the data is the linked bank-person information, allowing us to observe households bank connection. We exploit this information when analysing the natural experiment in Section 4.

2.4 Sample Restrictions and Summary Statistics

Our unit of analysis it the household, defined as individuals in a family relation living at the same address. We implement several sample selections to exclude observations where our electronic expenditure data may not accurately reflect consumption expenditure. First we exclude households how report self-employment income on the tax returns. For these individuals it difficult to separate between private and business related expenditures. We also exclude households-years when family members migrate, as we are less likely to cover consumption abroad. In addition, we remove households with very few electronic transactions. In particular we require that households in at least half of the months report a minimum of four debit card transaction and one wire transfer.

Our final sample consists of around 13 million household-year observations. Summary statistics is reported in Table 2.4.

	full sample	Decile of interest exposure									
		1	2	3	4	5	6	7	8	9	10
median age hh. head	45	65	56	42	47	48	48	45	42	39	36
median consumption/income	0.68	0.64	0.64	0.61	0.65	0.68	0.70	0.70	0.70	0.70	0.73
median debt/income	1.62	0.00	0.04	0.12	0.63	1.19	1.81	2.43	3.09	3.92	5.64
median deposits/income	0.21	2.22	0.64	0.10	0.16	0.19	0.18	0.15	0.13	0.11	0.11
median liquid assets/income	0.25	2.46	0.72	0.11	0.18	0.23	0.22	0.19	0.16	0.14	0.14
income after tax (2015 USD)	58,370	52,589	50,557	40,017	51,183	62,724	73,202	77,254	75,973	69,872	54,766
net worth (2015 USD)	91,724	408,193	203,854	3149	28,687	138,765	142,060	112,929	79,317	47,654	10,327
frac. single person hh.	0.42	0.49	0.48	0.56	0.45	0.38	0.32	0.29	0.30	0.35	0.49
frac. multi-person hh.	0.51	0.49	0.47	0.35	0.48	0.57	0.63	0.66	0.63	0.56	0.40
frac. with children in hh.	0.29	0.08	0.14	0.22	0.23	0.27	0.34	0.42	0.46	0.44	0.33
frac. head and spouse employed	0.61	0.34	0.48	0.51	0.58	0.62	0.66	0.70	0.73	0.74	0.70
frac. at least one employed	0.73	0.46	0.61	0.61	0.71	0.76	0.81	0.85	0.87	0.88	0.81
frac. homeowner	0.70	0.76	0.58	0.31	0.49	0.66	0.81	0.88	0.91	0.91	0.87
frac. higher education hh. head	0.32	0.29	0.26	0.18	0.26	0.32	0.34	0.34	0.36	0.39	0.43
observations (households-years)	13,342,546										

Notes. The table shows sample statistics for the final sample of households as well as by decile of interest exposure, defined as debt subtracted deposits and divided by income. Liquid assets are deposits, stocks and mutual funds. Higher education is education above high school level.

3 The Cash-flow Channel of Monetary Policy

When we estimate the effect of monetary policy on consumption, the fundamental identification problem we need to address is the endogeneity of monetary policy with respect to macroeconomic conditions. While monetary policy affects household consumption, it also reacts partly to other factors that drive consumption growth. For instance, a positive foreign demand shock increases output and consumption in the Norwegian economy, leading the central bank to increase the policy rate. Since the policy rate typically rises in response to an increase to shocks that drive up consumption, a projection of consumption growth on the change in interest rates would be biased towards zero.

In this section, we deal with this identification problem in two ways. In Section 3.1, we use time fixed effects to control for unobserved drivers of consumption growth that also affect the interest rate. This allows us to identify the differential effect of monetary policy on consumption between households that are expected to be exposed to a smaller or larger effect to changes in the interest rate based on their ex ante balance sheet composition. In Section 3.2, we instrument for changes to monetary policy using a new set of high-frequency instruments for Norway. This allows us to not only estimate the relative effect across households, but also the total effect of monetary policy on consumption. In addition, we test whether the consumption response is linear in ex ante exposure, measured as debt subtracted deposits. In section 3.3, we demonstrate that the previous estimates, under specific assumptions, can be understood as marginal propensities to consume out of the cash-flow change triggered by monetary policy. However, we also emphasize the need for caution in interpreting these effects, which leads us in Section 4 to estimate MPCs directly using exogenous variation in household interest expenses.

3.1 Identification With Fixed Effects

Empirical Strategy

We estimate how the consumption response to interest rate movements vary with households' ex ante exposure to monetary policy. Specifically, we estimate the local projection

(see Jordà, 2005)

$$\frac{c_{i,t+h} - c_{i,t-1}}{y_{i,\text{year}_t - 1}} = \beta^h \Delta r_t \times \text{EXP}_{i,\text{year}_t - 1} + \delta_i^h + \xi_t^h + \sum_{n=1}^{N} \gamma_{t,k}^{h,n} + X_{i,t}\alpha^h + \epsilon_{i,t}^h \tag{1}$$

separately for every horizon h, the number of months since the interest rate change Δr_t. The coefficient β^h measures how the effect of monetary policy on consumption varies with households' ex ante interest exposure $EXP_{i,t}$, defined as $\frac{b_{i,\text{year}_t-1} - d_{i,\text{year}_t-1}}{y_{i,\text{year}_t-1}}$, where both debt b and bank deposits d are measured at the end of the year previous to period t. As such, exposure is a pre-determined variable, but a single household can be part of different exposure groups at different points in time.[13]

Because our underlying consumption data is measured at a weekly frequency and the weeks do not exactly line up with calendar months, our dependent variable uses a three-month rolling average of consumption. Specifically, the variable $c_{i,t}$ is the average of the consumption of household i in the 12 weeks centered on month t.[14] Furthermore, y_{i,year_t-1} is the household's income in the previous year. We normalize the dependent variable by income in order to later compare our results to direct estimates of the marginal propensity to consume, as explained in section 3.3.[15]

We include multiple levels of fixed effects to control for endogeneity of monetary policy. In equation 1, the household fixed effects δ_i^h account for the differences in average consumption growth between households. The time fixed effects ξ_t^h account for variation in consumption growth between months that is common to all households. In addition, we include several fixed effects that are interactions between the time fixed effects and a pre-determined household observable, indexed by n. In our benchmark specification, we include $N = 9$ levels of these fixed effects: county of residence, age of household head, employment dummies for both household head and partner, industry of employment for

[13] We exclude observations with either values of the dependent variable or values of $EXP_{i,t}$ above the 99th and below the first percentiles.

[14] Since income is a yearly variable, we multiply the consumption variable by 400. Then, an increase in the dependent variable of 1 unit can be interpreted as an increase in consumption corresponding to 1 percent of income.

[15] This specification of the dependent variable is also the one used by Holm et al. (2021).

the household head, the number of people in the household, a homeowner dummy, and deciles for each of household income, wealth/income and liquid wealth/income.[16]

The purpose of including these fixed effects are twofold. First, by accounting for some of the variation in the household consumption growth, they reduce the variance of the error term $\epsilon_{i,t}^h$ and thereby potentially reduce the size of the standard errors. Second, and most importantly, they can remove the component of the error term that is correlated with the interest rate change Δr_t. The OLS estimator of β^h is consistent if changes in the policy rate are correlated with unobserved drivers of consumption growth that are common to households within the groups that the fixed effects account for. For instance, the central bank could partly set the interest rate based on employment growth that varies between regions (accounted for by time-county fixed effects), or based on shocks to income growth that affect consumption equally within age cohorts (accounted for by time-age fixed effects). Furthermore, the time-month fixed effects account for variation in the interest rate that that correlates with the calendar month.

We also include a vector of household control variables $X_{i,t}$. This vector includes the exposure variable and the lagged dependent variable.[17] We also include 12 month dummies as well as interacting these dummies with the exposure variable. Since the regression is by horizon h, the effects of all controls on the dependent variable are allowed to vary flexibly by horizon. The error term $\epsilon_{i,t}^h$ is specific to a household, month and horizon. We cluster standard errors by time and household.

We use the 3 month money market rate (the Norwegian Interbank Offered Rate, NIBOR) as our measure of r_t. NIBOR is the reference rate used by Norwegian banks when setting their lending and deposit rates. It also closely follows the main policy rate set by Norges Bank. Figure 9 shows that all these rates track each other over time, but

[16] It is computationally too costly to run regression 1 on the full panel of Norwegian households for all years and months. For that reason, we have estimated the regression on a 15% random sample of households. We sample randomly from the full list of identifiers of households that are in our dataset in at least one year between 2006 and 2018. If a household is sampled, it stays in our sample for all years it is present. We have done the sampling two times and checked that the results do not change across samples.

[17] All balance sheet variables used for controls are measured at the end of the year before period t. The lagged dependent variable is defined as $\frac{c_{i,t-2}-c_{i,t-5}}{y_{i,\text{year}_t-1}}$.

with different levels.[18]

Results

Figure 5 plots the estimates for β^h in regression 1 by horizon (month) h. It shows the differential response to a 1 percentage point increase in the money market rate between households with high and low interest exposure. In figure 5a, we show the estimates of the specification with only time fixed effects. In figure 5b, we include the full set of fixed effects. The figures show that the two specifications give similar results.

The interpretation of the point estimate -0.077 for the first month is that the median household, which has an exposure equal to 1.25, reduces consumption by 0.097 (0.077 times 1.25) percentage points more, as a fraction of income, than a household with zero exposure if the interest rate increases by 1 percentage points. Our estimates indicate that the differential response increases after around 5 months and reaches a peak in month 8. In month 6, the median household reduces consumption by 0.29 percentage points more (0.23 times 1.25) than not-exposed households. At this horizon, households at the 90th percentile reduces consumption by 1.01 percentage points more than not-exposed household. Except for the first few months, the estimates are significant at the 95 percent level.

Combining our coefficient estimates with the median consumption-to-income ratio of 0.68 (see table 2.4), a back-of-the-envelope calculation indicates that consumption falls by approximately 0.43 percentage points more (0.23 divided by 0.68) after 6 months for the median household, and 1.49 percentage points more for a household at the 90th percentile, compared to the not-exposed households. Our estimates are slightly smaller than the comparable results for Swedish data from Flodén et al. (2020).

[18]While the lending and deposit rates are more directly related to the cash-flow changes experienced by households, these rates are only available at a monthly frequency starting in December 2013.

Figure 5: Consumption response to interest rate increases.

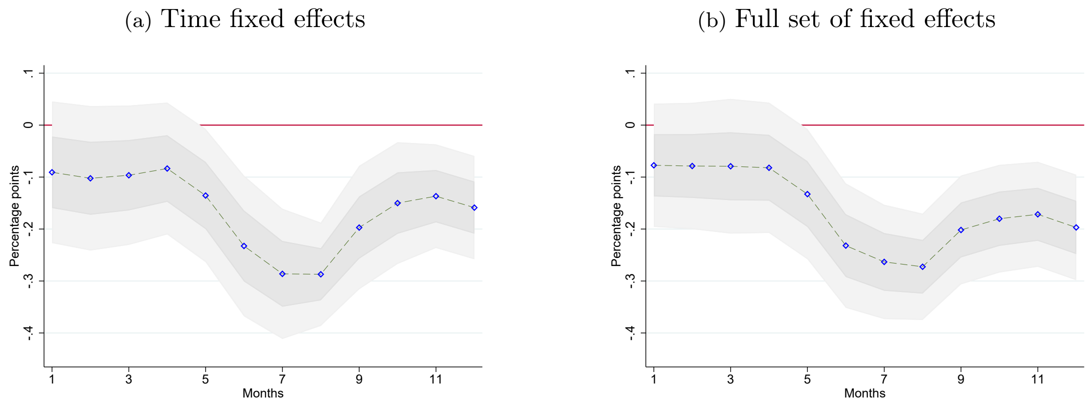

Notes. The figures show the estimates of coefficient β^h from regression 1 for horizons (months) $1-12$. This coefficient can be interpreted as the additional response in consumption to a 1 percentage point higher interest rate when ex ante interest exposure is 100 percent higher. The figure on the left-hand side shows the estimates for the regression that only includes time fixed effects, while the right-hand side figure shows the estimates for the regression with a full set of time-household fixed effects.

3.2 Identification With Monetary Policy Instruments

Empirical Strategy

We now order households into quantiles of their interest exposure and estimate how much the response of consumption to a change in the interest rate varies over time and across the quantiles. To deal with the endogeneity of interest rates, we construct a set of high-frequency monetary policy instruments. Using the instruments allow us to identify both the heterogneity in consumption responses between the groups of households and the level of the effect, since we no longer need to include time fixed effects.

Specifically, our empirical strategy is to estimate the local projection

$$\frac{c_{i,t+h} - c_{i,t-1}}{y_{i,\text{year}_t - 1}} = \beta_g^h \Delta r_t + \delta_i^h + X_{i,t}\alpha_g^h + \epsilon_{i,t}^h \quad \forall i \in \mathcal{I}_g \tag{2}$$

for every horizon h and every group of households $\mathcal{I}_g$ separately. We include in the vector $X_{i,t}$ household income in the previous year, total wealth relative to income, liquid wealth relative to income, and the lagged dependent variable. We also include 12 month dummies and a second order polynomial in time. Since the regression is by horizon h, the effects of all controls on the dependent variable are allowed to vary flexibly by horizon.

The coefficient β_g^h now identifies the average effect within group $\mathcal{I}_g$ of a one percentage point increase in the interest rate on consumption after h months, measured in units of income. The construction of the groups is explained below.

Identification of Monetary Policy Instruments

We construct a new set of monetary policy instruments for Norway based on high-frequency changes in market expectations on announcement days. When Δr_t is instrumented, the regression 2 is a local projection instrumental variables (LP-IV) regression (Stock and Watson, 2018).

The construction of our monetary policy instruments follows a two-step procedure. First, following a large literature (see Gürkaynak et al. (2005), Gertler and Karadi (2015), and Jarociński and Karadi (2020), among others), we extract the surprise element of monetary policy announcements from high-frequency changes in market-based expectations on announcement days. The expectations are based on the pricing of four forward rate agreements (FRAs), which gives us a total of four instruments. These instruments reflect the change in the market's expectations of the 3 month money market rate in a 30 minute window around the announcement time, from 10 minutes before the announcement to 20 minutes after the announcement.[19] The shorter the window, the less we expect the expectations to be affected by other macroeconomic conditions that might also be correlated with consumption growth. At the same time, it is important to allow enough time for the market to process the monetary policy announcement.

Second, we adjust our instruments to account for the information effect of monetary policy decisions. As argued by Blinder et al. (2008), among others, central banks provide the public with information about the macroeconomic outlook as well as the outlook for their own decisions. Some of this information is transmitted on announcement days in the form of central bank forecasts for macroeconomic variables. This might invalidate

[19]The same FRA contracts have previously been used by Brubakk et al. (2022) to generate monetary policy instruments for Norway.

the exclusion restriction of the instrument, for instance when a policy hike is a signal that the macroeconomic outlook is stronger than previously assumed by agents in the economy. We follow Miranda-Agrippino and Ricco (2021) in adjusting the market-based surprises for this information component. First, we project the surprises on a set of forecasts and forecast revisions for inflation and output growth prepared by Norges Bank ahead of each meeting of the Monetary Policy Committee. These forecasts might contain information about the state of the economy that have not already been internalized by agents in the economy. Second, we extract the residuals from this regression and project them on their own lags, thereby removing a potential autoregressive component of the surprises that can be due to slow absorption of information among market participants (Miranda-Agrippino and Ricco, 2021). The residuals from this second regression are our final instruments. See appendix B for a detailed description of how the instruments are constructed.

Interest exposure

We first order the households into 20 *ventiles* – 5 percent groups – based on the interest exposure variable $EXP_{i,t}$ used in section 3.1. We then run regression 2 separately for each group and each horizon.

In table 2.4 we summarize the sample characteristics of each of the exposure groups. The median interest exposure increases from -2.0 in the first decile (bottom two ventiles) to 5.4 in the top decile (top two ventiles). Households with negative exposure tend to be older and have higher net worth but lower income than the median household. Households with high levels of exposure to the interest rate changes have lower net worth, but are more likely to be homeowners and have children in the household.

Figures 13 and 14 show the effect of a one percentage point increase in the interest rate on consumption by exposure. Each subfigure contains a particular horizon, from month 1 (impact) to month 12.[20] We focus here on the first year and describe the estimates for

[20]The estimated consumption responses are statistically significant at the 68% level for higher exposure groups, but not at the 95% level.

the second year below. Figure 6a summarizes the estimates in a single figure containing four of the horizons. To aid legibility, we add flexible splines fitted to the estimates for each of the horizons separately.

There are two things to note from the results. First, the point estimates indicate that there is no response in consumption on impact (month 1) for any of the groups, but by month 5 the consumption response is increasing in exposure, with a response of around -0.7 percent of income at the top of the distribution. This number rises to around -2 after one year.

Second, figure 6a shows that there is a non-linearity in the estimated responses across the size of the exposure. Households with negative exposure do not appear to respond to the interest rate changes at any horizon, while the consumption response is close to linear in exposure for households who have more debt than deposits. As we see in table 2.4, households with negative exposure typically have very little debt. Instead, decreasing exposure is associated mostly with increasing deposits for these households. For households with positive exposure, increasing exposure is mostly associated with increasing levels of debt, as these households on average hold little deposits.

Debt and deposits

To further investigate the separate role of deposits and debt, we now order the households by their debt-to-income (DTI) ratio and rerun regression 2. The results for each horizon are shown in figures 15 and 16, while figure 6b shows a subset of horizons. We see that the consumption response is close to linear across the whole distribution of debt-to-income. Hence, it appears that there is no cash-flow effect on consumption through households' deposits, while there is a clear effect through debt.

This difference can be due either to a smaller, or slower, pass-through of changes in the policy rate to deposit rates than to lending rates, or it can be due to a smaller marginal propensity to consume out of the interest earnings on deposits. To investigate this, we estimate pass-through regressions by projecting separately the average monthly lending rate and the average monthly deposit rate at various horizons on the initial change in the

Figure 6: Consumption response to interest rate increases.

(a) By exposure

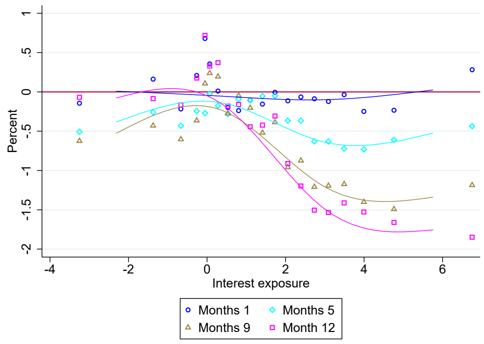

(b) By debt/income

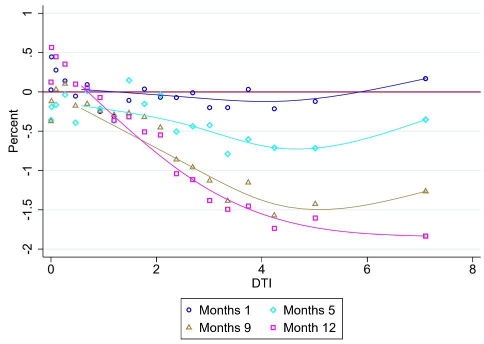

Notes. The figure shows the results from regression 2 by time horizon and quantiles (20 ventiles) of households based on debt-to-income (DTI). Month 0 is the month of the instrumented interest rate change. Each dot is a separate estimate of coefficient β_g^h of the regression for a particular ventile g and horizon h. In figure 6a, households are ordered by interest exposure, while in figure 6b they are ordered by debt-to-income. The median value of the grouping variable for each of the ventiles is shown on the horizontal axis. The vertical axis shows the estimated response of consumption, relative to its level in the three months before the change in the interest rate, in percent of income. To ease exposition, we fit a spline to the dots separately by horizon.

money market rates.[21] The results are shown in figure 12. The lending and deposit rates exhibit very similar dynamics following an increase in the money market rate. They both peak at an increase of around 1 percentage point after one year. While not conclusive, this points in the direction of a different consumption response to a change in the deposit rate than to the same change in the lending rate.

The second year

Figures 17 and 18 show the responses for each of the horizons within the second year after the interest rate change. These results indicate that the cash-flow effects are smaller in the second year than in the first year. This is confirmed by figure 19, which shows the average of the responses over the first year for the bottom 20 percent of households based on exposure, the middle 20 percent and the top 20 percent. The figure shows that

[21] The regression we estimate is $r_{t+h}^{hh} = \alpha_0^h + \alpha_1^h r_{t-1}^{hh} \alpha_2^h \Delta r_t + \varepsilon_t^h$, where r_{t+h}^{hh} is either the lending or deposit rate at horizon $t+h$ and Δr_t is the change in the money market rate between month $t-1$ and month t. We use the interest rates on all outstanding repayment loans and on deposits for the household sector, both from a sample of banks and mortgage companies and made available by Statistics Norway. These rates are only available starting in December 2013. Hence the pass-through estimates are based on less than half of the years in our sample, and the results should be interpreted with caution.

Figure 7: Cash-flow effects of monetary policy

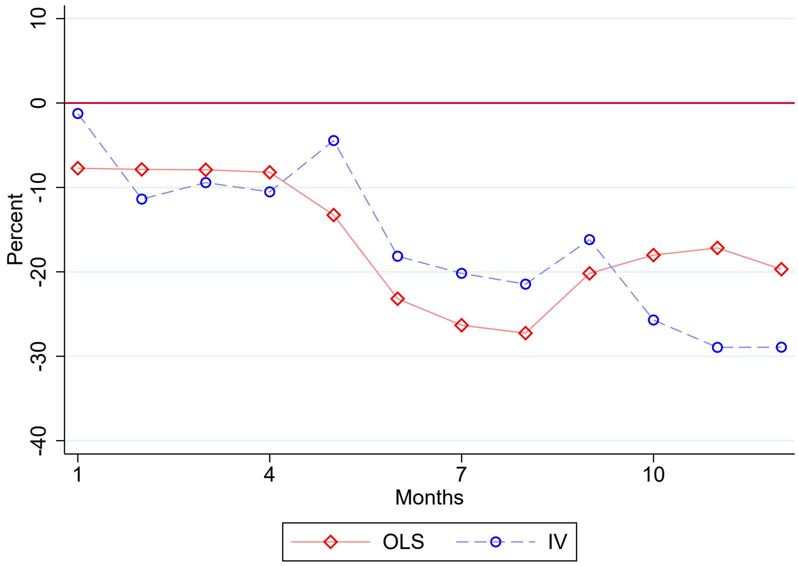

Notes. The figure shows implied marginal propensity to consume (MPC) estimates from regression 1 (OLS) and 2 (IV), by horizon (month) after a change in the interest rate. The interpretation of these estimates as MPCs requires the assumption that the cash-flow change resulting from the interest rate change Δr_t is given by $(b_{i,\text{year}_t-1} - d_{i,\text{year}_t-1}) \times \Delta r_t$, where b_{i,year_t-1} and d_{i,year_t-1} are the debt and deposits, respectively, of household i in the year prior to the interest rate change.

the consumption response of high exposure households falls after around one year and might in fact be reversed.

To interpret these results, we should take into account how the interest rate moves over time after an initial increase. Figure 11 shows the estimates from a local projection of the money market rate on an initial increase in the same rate, instrumented with our monetary policy instruments. As in regression 2, we include month dummies and a second order polynominal in time. The interest rate jumps by 1 percentage point on impact and then continues to increase gradually until it peaks at around 1.5 percentage points after 9 months. The effect of the initial shock disappears after one year.

3.3 Cash Flow Effects and the Marginal Propensity to Consume

In order to interpret the estimates from section 3.3 and 3.2 as *cash-flow effects* of monetary policy, the exposure variable multiplied by the interest rate change must be a measure of the cash-flow change induced by the interest rate. If that is the case, we see directly from inspecting regression 1 that the coefficient β^h is a measure of the marginal propensity to consume (MPC) out of net interest expenses. To get an equivalent measure from our IV results, we estimate a linear regression on the individual point estimates for a particular horizon. The slope of the line is an estimate of the MPC. In figure 7 we see that these back-of-the-envelope MPC estimates are similar in size and increase at the same rate over the horizons. The implied MPCs out of net interest expenses are around 10 percent in the first half of the year following the initial interest rate change, and they increase to around $20 - 30$ percent at the end of the year. On average through the year, both methods give an MPC estimate of 16 percent.

At this point it is worth highlighting several caveats with this MPC interpretation by inspecting the right-hand side variable in regression 1, $\frac{b_{i,\text{year}_t-1} - d_{i,\text{year}_t-1}}{y_{i,\text{year}_t-1}} \times \Delta r_t$. First, households' are not uniformly exposed to the changes in the money market rate measured by Δr_t, both because the pass-through of money market rates to lending and deposit rates might not be uniform and, crucially, because the pass-through can differ across households. For instance, households with a high debt-to-income ratio might be more likely to negotiate with their existing bank or switch to a different bank when interest rates rise compared to households with a low debt-to-income ratio. Second, both the exposure variable and the interest rate change are measured at an earlier point in time than the consumption change. Figure 11 shows that the interest rate on average continues to increase for several months after an initial contractionary shock. Second, households' balance sheet position at the end of the previous year might not accurately capture the debt and deposits that are subject to changes in interest rates within the current year. Lastly, as discussed in section 3.2, the propensity to consume out of the interest earned on deposits might not be the same as the propensity to consume out of interest expenses.

It follows from this discussion that in order to more accurately measure the MPC out of interest expenses, we need households' net interest payments to be both (a) directly observed over the same period as consumption and (b) subject to an exogenous shock that only affects either deposits or debt. In the next section we show how these conditions arise in a natural experiment setting.

4 Household-specific Interest Rates: A Natural Experiment

4.1 Institutional Background

Since the early 1900s, Norway has operated a government-owned lending institution called Statens Pensjonskasse (SPK), which provides floating-rate-only mortgages to public sector workers at favorable rates.[22] Historically, the mortgage rate, which is universal and non-negotiable, was a markup over 3-month treasury bills. SPK mortgages were subject to a borrowing limit of 200,000 USD (2023 exchange rate) per individual and a loan-to-value cap of 80 percent.[23]

In contrast, following the deregulation of the Norwegian banking sector in the 1980s, mortgage rates in conventional banks were determined as a markup over the 3-month money market rate. Unlike the markup in conventional banks, which is an equilibrium object influenced by demand and bank competition, the markup in SPK was a political decision made by the government annually as part of the National Budget. Historically, the SPK markup has been fixed at 50 basis points, resulting in the SPK borrowing rate being the best mortgage rate offer in the market and about one percentage point lower than the average bank mortgage rate.[24]

In the aftermath of the Great Financial Crisis (GFC), a structural shift materialized, causing the SPK spread to increase well above historical levels as shown in Figure 8.[25] In the fall of 2013, as part of the 2014 National Budget process, the government decided to

[22]In addition to offering mortgages to its members, SPK also operates an occupational pension scheme. It is not, however, a deposit taking institution.

[23]For some households the borrowing limit has been binding, leading them to combine SPK debt with conventional bank mortgages.

[24]Because of the way SPK sets interest rates, the effective fixed-rate period was approximately 4 months, whereas in conventional banks, the corresponding lock-in period was 6 weeks. During periods of rapidly declining rates, this meant that the SPK spread could temporarily become negative.

[25]The exact reason for this increase is unclear, but it likely stemmed from a combination of factors including a widened spread between treasury bills and money market rates, as well as reduced competitive pressure in the banking sector.

address this by raising the SPK markup by 75 basis points, effective from March 1st, 2014. The policy shift was not communicated to the public before the National Budget was published on October 8, 2013. The increased markup resulted in a rate hike which was communicated to SPK customers in January 2014. This action was repeated in the 2015 budget process, resulting in a further 50 basis point increase in March 2015. Subsequently, in the 2016 National Budget, the government opted to introduce a completely new pricing formula, which remains in effect to this day. Under this formula, the SPK rate is set at 15 basis points below the average of the five best mortgage offers available in the market at any given time.

As a result of the series of policy reforms initiated in 2014, there was a substantial and rapid decrease in the interest differential, ultimately restoring the spread to its pre-GFC level by 2016, as shown in Figure 8. Our empirical strategy, outlined below, leverages the 2014 policy change as a starting point and compares the trajectory of interest expenses and consumption between SPK and non-SPK customers.[26]

[26]Since the interest rate differential is positive and time-varying even before 2013 (see figure 8), there will likely be differences in consumption dynamics between the two groups that are driven by these mortgage rate movements even in the earlier years of our sample. We choose to focus on the years immediately after 2013 since the drivers of the interest rate differential in those years are clearly understood and exogenous to unobserved drivers of consumption.

Figure 8: Interest Rate Differential

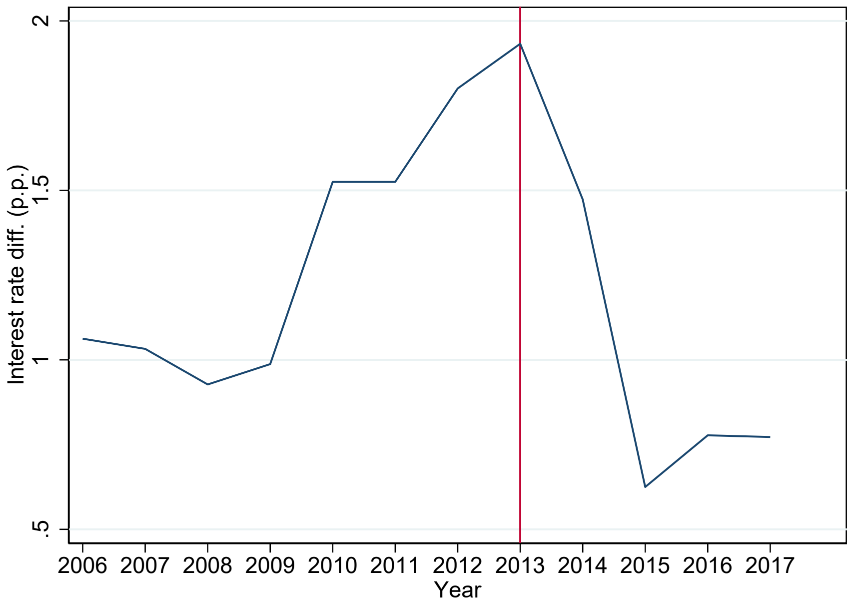

Notes. Source: Statistics Norway. This figure plots the percentage point difference in loan rates between all household debt and SPK mortgages. The rates are loan-weighted average of existing household level debt.

4.2 Sample Definitions and Descriptive Statistics

We conduct the analysis in this section at the yearly frequency in order to directly observe changes in household interest expenses. We label households as SPK or non-SPK households based on whether or not the household had an outstanding stock of debt in SPK on December 31, 2013. We identify $66,230$ households as SPK customers at the end of 2013. To improve statistical power we remove from our sample households labeled as SPK if (i) they are not customers of SPK either the year before or the year after 2013, and (ii) they do not display a stable debt repayment trajectory in the two year period from December 31, 2012 to December 31, 2014.[27] The purpose is to remove households with large changes in interest payments to SPK between 2013 and 2014 driven by events unrelated to the rate hike in SPK, for instance housing market transactions or

[27] A stable trajectory is defined as SPK debt declining each year, but with an annual re-payment below 10 percent of the beginning-of-year stock of debt.

labor market moves in and out of the public sector.[28] After imposing restriction (i) and (ii) we are left with 22, 699 SPK households. Of the households we lose, the majority (90 percent) is due to extensive margin debt movements in and out of SPK over the 2 year period. We impose similar stable debt restrictions for the control group consisting of non-SPK households, based on the debt trajectory in their main bank[29]

In addition to the stable debt trajectories, we further impose the same type of sample restrictions as in Section 3.

4.3 Empirical Strategy

We estimate the regression

$$Y_{i,t} - \overline{Y}_i = \alpha_Y^t + \beta_Y^t D_i + X_i \gamma_Y^t + \epsilon_{i,t}, \tag{3}$$

where $Y_{i,t}$ is the outcome variable of interest (consumption or gross interest expenses) and $\overline{Y}_i$ is the value of this variable for household i in the reference year 2013, while $D_i \in \{0, 1\}$ is a dummy variable indicating whether household i is in the treated (SPK) group. We run this regression separately for every year t, including the same set of treated and non-treated households in every regression.

The constant α_Y^t captures time variation in average consumption. The vector X_i includes a large set of pre-determined household control variables that might account for variation in household consumption and interest payments. In our benchmark specification, we include 10 year age groups of the household head, dummy variables for the number of adults and the number of children in the household, the maximum education level for the household head and spouse (four levels), and dummy variables for the decile of the level of debt, the level of income and the level of consumption relative to income in 2012.

[28] A potential concern arising from this restriction could be that a large share of SPK clients switched to other banks because of the 2014 reform. However, the SPK offer remained the best possible mortgage rate also after the reform, which mitigates this concern. In addition, the exit rate from SPK is stable at 5 percent each year between 2010-2014

[29] The main bank of non-SPK households is defined as the bank in which the household has its largest debt share.

Table 2: MPC out of interest expenses.

	Year		
	2014	2015	2016
interest exp.	912^{***}	1985^{***}	2154^{***}
	(8.9)	(16.1)	(20.6)
consumption	-278^{**}	-825^{***}	-835^{***}
	(86.9)	(118)	(137)
MPC	30.5^{**}	41.5^{***}	38.8^{***}
	(11.1)	(6.62)	(7.48)
Observations	$289,182$	$289,168$	$273,786$

Notes. The table shows the estimated coefficient β^t in regression 3 for each of the years $t \in \{2014, 2015, 2016\}$ and for both interest expenses and consumption. The coefficient estimates measure the effect of treatment (SPK households vs. non-SPK households) on the outcome variables. The column also shows our estimate for the MPC out of interest expenses, derived from the former estimates. Standard errors (in parentheses) are bootstrapped using 200 draws.

We estimate regression 3 separately by year and by outcome variables interest expenses and consumption, using 2013 as the reference year. β^t_Y then measures the difference in the outcome variable in year t between SPK and non-SPK households, controlling for other observables. Our estimate of the MPC out of interest expenses in year t is given by $-\frac{\beta^t_{\text{consumption}}}{\beta^t_{\text{interest exp.}}}$. Standard errors of the MPC estimates are bootstrapped with 200 draws.

4.4 Results

The estimates for regression 3 are shown in table 2. All else equal, the average interest expenses of an SPK household increase by USD 912 relative to a non-SPK household in 2014, while consumption falls by USD 278 more in the former group. As a result, our estimated MPC out of interest expenses is 30.5 percent in 2014. All our estimates are significant at the 95 percent level.

In the second and third columns of table 2 we report the same estimates for the years 2015 and 2016 relative to the reference year, the effect of treatment on interest expenses is larger than in 2014, reflecting the fall in the interest rate differential between SPK and other banks throughout 2015 shown in figure 8. We estimated marginal propensities to consume out of these interest expenses are larger in 2015 and 2016.

Our estimates of the MPC out of interest expenses are similar to, but slightly larger than those reported in figure 7 for the other two identification methods used in this paper, in particular for the years 2015 and 2016. One reason for this might be that the movements in net interest expenses induced by the SPK reform were in fact – and were most likely expected to be – more permanent that those induced by typical movements in the central bank policy rate. It is also worth noting that MPC estimates in the range $30-40\%$ are similar to, but at the low end of, those typically reported out of other exogenous income shocks, such as unemployment shocks (Fagereng et al., 2024), lottery prizes (Fagereng et al., 2020) and tax rebates (Parker et al. (2013); Jappelli and Pistaferri (2014); Misra and Surico (2014)).

5 Conclusion

We investigate how monetary policy affects consumption through the interest exposure of households. By employing several identification methods on a new, comphrehensive and high-frequency dataset of directly measured consumption expenditures at the household level, we estimate the response of consumption to an unexpected monetary policy shock along the distribution of net interest exposure and debt, as well as the marginal propensity to consume out of changes in interest payments. Our results indicate that there is a substantial cash-flow channel of monetary policy operating over a relatively short horizon, and we find MPCs out of exogenous movements in interest expenses that are similar in size to those estimated out of other types of shocks to disposable income.

References

Aastveit, K. A., K. R. Gerdrup, and A. S. Jore (2011): "Short-term forecasting of GDP and inflation in real-time: Norges Bank's system for averaging models," Discussion Paper 9, Norges Bank.

Aguiar, M. A., M. Bils, and C. Boar (2020): "Who are the Hand-to-Mouth?," Discussion paper, National Bureau of Economic Research.

Alpanda, S., E. Granziera, and S. Zubairy (2021): "State dependence of monetary policy across business, credit and interest rate cycles," *European Economic Review*, 140, 103936.

Andersen, A. L., A. S. Jensen, N. Johannesen, C. T. Kreiner, S. Leth-Petersen, and A. Sheridan (2021): "How do households respond to job loss? Lessons from multiple high-frequency data sets," .

Auclert, A. (2019): "Monetary policy and the redistribution channel," *American Economic Review*, 109(6), 2333–2367.

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen (2008): "Central bank communication and monetary policy: A survey of theory and evidence," *Journal of economic literature*, 46(4), 910–945.

Boehm, J., E. Fize, and X. Jaravel (2023): "Five Facts about MPCs: Evidence from a Randomized Experiment," .

Brubakk, L., S. Ter Ellen, Ø. Robstad, and H. Xu (2022): "The macroeconomic effects of forward communication," *Journal of International Money and Finance*, 120, 102536.

Buda, G., V. M. Carvalho, G. Corsetti, J. B. Duarte, S. Hansen, Á. Ortiz, T. Rodrigo, and J. V. Rodríguez Mora (2023): "Short and Variable lags," .

Carroll, C., J. Slacalek, K. Tokuoka, and M. N. White (2017): "The distribution of wealth and the marginal propensity to consume," *Quantitative Economics*, 8(3), 977–1020.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of political Economy*, 113(1), 1–45.

Clarida, R., J. Gali, and M. Gertler (1999): "The science of monetary policy: a new Keynesian perspective," *Journal of economic literature*, 37(4), 1661–1707.

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

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao (2017): "Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging," *American Economic Review*, 107(11), 3550–3588.

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

Fagereng, A., M. B. Holm, and G. J. Natvik (2021): "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 13(4), 1–54.

Fagereng, A., H. Onshuus, and K. N. Torstensen (2024): "The consumption expenditure response to unemployment: Evidence from Norwegian households," *Journal of Monetary Economics*, p. 103578.

Flodén, M., M. Kilström, J. Sigurdsson, and R. Vestman (2020): "Household Debt and Monetary Policy: Revealing the Cash-Flow Channel," *The Economic Journal*, ueaa135.

Ganong, P., D. Jones, P. J. Noel, F. E. Greig, D. Farrell, and C. Wheat (2020): "Wealth, race, and consumption smoothing of typical income shocks," Discussion paper, National Bureau of Economic Research.

Gelman, M. (2022): "The Self-Constrained Hand-to-Mouth," *Review of Economics and Statistics*, 104(5), 1096–1109.

Gerard, F., and J. Naritomi (2021): "Job displacement insurance and (the lack of) consumption-smoothing," *American Economic Review*, 111(3), 899–942.

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

Gürkaynak, R. S., B. Sack, and E. 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.

Gürkaynak, R. S., B. P. Sack, and E. T. Swanson (2004): "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements," *The Response of Asset Prices to Monetary Policy Actions and Statements (November 2004).*

Hamilton, S., G. Liu, and T. Sainsbury (2023): "Early pension withdrawal as stimulus," *Available at SSRN 4389699.*

Holm, M. B., P. Paul, and A. Tischbirek (2021): "The transmission of monetary policy under the microscope," *Journal of Political Economy*, 129(10), 2861–2904.

International Monetary Fund (2024): "World Economic Outlook – Steady but Slow: Resilience amid Divergence," .

Jacobson, M., C. Matthes, and T. B. Walker (2022): "Inflation Measured Every Day Keeps Adverse Responses Away: Temporal Aggregation and Monetary Policy Transmission," *Finance and Economics Discussion Series*, (2022-054).

Jappelli, T., and L. Pistaferri (2014): "Fiscal policy and MPC heterogeneity," *American Economic Journal: Macroeconomics*, 6(4), 107–136.

Jarociński, M., and P. Karadi (2020): "Deconstructing monetary policy surprises—the role of information shocks," *American Economic Journal: Macroeconomics*, 12(2), 1–43.

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

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

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary policy according to HANK," *American Economic Review*, 108(3), 697–743.

KUTTNER, K. 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, G., H. HUGHSON, G. KAPLAN, ET AL. (2016): *The household cash flow channel of monetary policy*. Reserve Bank of Australia.

LEWIS, D. J., D. MELCANGI, AND L. PILOSSOPH (2019): "Latent heterogeneity in the marginal propensity to consume," *FRB of New York Staff Report*, (902).

MIAN, A., K. RAO, AND A. SUFI (2013): "Household balance sheets, consumption, and the economic slump," *The Quarterly Journal of Economics*, 128(4), 1687–1726.

MIRANDA-AGRIPPINO, S., AND G. RICCO (2021): "The transmission of monetary policy shocks," *American Economic Journal: Macroeconomics*, 13(3), 74–107.

MISRA, K., AND P. SURICO (2014): "Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs," *American Economic Journal: Macroeconomics*, 6(4), 84–106.

NORGES BANK (2023a): "Financial Infrastructure Report," Discussion paper, Oslo, Norway.

——— (2023b): "Retail payment services 2022," *Norges Bank Papers*, (1).

PARKER, J. A. (2017): "Why don't households smooth consumption? Evidence from a $25 million experiment," *American Economic Journal: Macroeconomics*, 9(4), 153–183.

PARKER, J. A., N. S. SOULELES, D. S. JOHNSON, AND R. MCCLELLAND (2013): "Consumer spending and the economic stimulus payments of 2008," *American Economic Review*, 103(6), 2530–2553.

PICA, S. (2023): "Housing Markets and the Heterogeneous Effects of Monetary Policy Across the Euro Area," .

RAMEY, V. A. (2016): "Macroeconomic shocks and their propagation," *Handbook of macroeconomics*, 2, 71–162.

STOCK, J. H., AND M. W. WATSON (2018): "Identification and estimation of dynamic causal effects in macroeconomics using external instruments," *The Economic Journal*, 128(610), 917–948.

WONG, A., ET AL. (2019): "Refinancing and the transmission of monetary policy to consumption," *Unpublished manuscript*, 20(4).

A Credit card imputation

Payments to banks (category 13) include debt service payments, in particular mortgage payments. We impute and remove these payments by splitting category 13 invoice payments in two groups. The grouping is based on the whether a payment is likely to be consumption (e.g. credit card) or investment (e.g. mortgage) related, and we clean the data by removing transactions classified as the latter. The imputation is based on the assumption that service of investment related debt is relatively stable over time, in contrast to other bank payments (e.g. credit card bills). The algorithm consist of 7 steps performed sequentially per individual. In each step, the payment is classified as imputed debt service if the condition is met and the transaction amount is above 2500 NOK (2018 CPI adjusted).

1. Local stable payment: Compute the moving coefficient of variation (standard deviation over mean) over a rolling window of 8 payments within a zip-code, dropping the largest and smallest payment.[30] for payments with multiple transactions within a week we first compute the average payment by dividing with number of transactions. 7.2 percent of invoice observations in category 13 involve multiple transactions. The payment is classified as imputed debt service if the coefficient of variation is below 0.12.

2. Typical payment within a year: Re-classify the remaining 25 percent of transactions as imputed debt service if the first step classifies more than 75 percent of transaction within the year as imputed debt service.

3. First and last year adjustment (I): Classify all payments in the individual's first (last) year of observation as imputed debt service if all payments in the following (preceding) year are classified as imputed debt service. This adjustment is done because the moving coefficient of variation is missing for the first and last 4 observations.

4. First and last year adjustment (II): Assume that the first (last) 4 observations are imputed debt service payments if the 5th (T-4) observation is a mortgage payment.

5. Temporary non-mortgage: Re-classify as imputed debt service if more than 75 percent of observations in a window of +/- 2 observations are classified as mortgage.

6. Typical payment in full sample: Re-classify as imputed debt service if more than 70 percent of all observations are imputed debt service.

7. Very large payment: Re-classify as imputed debt service is transaction amount is above 100,000 NOK (2018 CPI adjusted)

[30] For payments with multiple transactions within a week we first compute the average payment by dividing with number of transactions. 7.2 percent of invoice observations in category 13 involve multiple transactions.

Some payments in category 13 we know are related to debt service in a certain mortgage bank and student debt bank, and we separate out these two before applying the algorithm.[31] We can precisely identify these payments because in certain zip codes there are only one bank. For two of these unique banks, individuals' bank service is solely related to either mortgage or student debt We nevertheless apply the algorithm on payments to these banks to check its performance. For the mortgage bank, our algorithm classifies 82.5 percent of the payments as debt service. For the student debt bank, it classifies 91.2 percent of the payments.[32]

[31] We can precisely identify these payments because in certain zip codes there are only one bank. For two of these unique banks, individuals' bank service is solely related to either mortgage or student debt.

[32] For the student debt bank, we do not apply the requirement that the transaction amount must be above 2500 NOK for it to be classified as imputed debt service.

B Construction of monetary policy instruments

We construct four separate monetary policy instruments based on the first four forward rate agreements (FRA) based on the 3 month NIBOR rate. An FRA specifies a future interest rate and is used by market participants to hedge against interest rate risk. At a pre-determined expiration date, the seller of the contract pays the buyer the difference between the 3 month NIBOR rate at that date and the rate specified in the contract if that difference is positive, and vice versa. Hence, the FRA rate reflects the market's expectation of the 3 month money market rate at the expiration date plus a potential forward premium. The first FRA uses the upcoming International Money Market (IMM) date as its expiration date, the second FRA uses the next IMM date, and so on. Since there are four IMM dates in a year, the instruments reflect expectations of the money market rate from within the first quarter after a monetary policy announcement until around one year after the announcement. We construct our instruments for the time period from the monetary policy announcement in July 2004 until the last announcement in 2018.[33] The data is extracted from the 1 minute Thomson Reuters Tick History database maintained by Refinitiv.

Let FRA_d^i be the change in expectations associated with FRA contract i on announcement day number d ($d = 1$ is the first announcement day in our sample, $d = 2$ is the second one, and so on). Following Miranda-Agrippino and Ricco (2021), we run the regression

$$FRA_d^i = \alpha^i + \sum_{j=0}^{3} F_d^{NB} x_{q+j} \theta_j^i + \sum_{j=0}^{2} [F_d^{NB} x_{q+j} - F_{d-1}^{NB} x_{q+j}] \gamma_j^i + \overline{MPI}_d^i, \tag{4}$$

where $F_d^{cb} x_{q+j}$ is the forecast prepared by Norges Bank before announcement day d for the vector of variables x at horizon $q+j$. $F_d^{NB} x_{q+j} - F_{d-1}^{NB} x_{q+j}$ is the associated forecast revision. Here j denotes the number of quarters from the quarter q in which the forecast is prepared. Here $\overline{MPI}_d^i$ is the residual of the regression that uses FRA number i.

Over our sample period, Norges Bank has prepared and made public two types of forecasts. The first are the official forecasts published along with the triannual (until 2012) and quarterly (from 2013) Monetary Policy Reports. These forecasts are not suited to our needs, since they are only available for around half of the monetary policy meetings in our sample. Instead, we rely on a set of model-based forecasts produced by Norges Bank's system for averaging models (SAM). SAM forecasts are typically updated multiple times each quarter, and they are used as starting points to generate the official forecasts for the Monetary Policy Reports. They have typically also been published along with the official forecasts (see Aastveit et al. (2011) for a detailed description). We use the SAM forecasts for the quarterly GDP growth rate and the quarterly core inflation rate (CPI-ATE) in regression 4.

We now aggregate the residuals $\overline{MPI}_d^i$ from regression 4 to the monthly level by

[33] While the FRA data is available from 2001, the macroeconomic forecasts are only available starting in January 2004, and hence the first forecast revisions are available in the second quarter of 2004.

summing over meeting days within a month (there are very few months in which there are multiple meetings). We then project the monthly residuals $\overline{MPI}_t^i$ on twelve of their own lags by running the regression

$$\overline{MPI}_t^i = \phi_0^i + \sum_{j=1}^{12} \phi_j^i \overline{MPI}_{t-j}^i + MPI_t^i. \tag{5}$$

As explained by Miranda-Agrippino and Ricco (2021), this step accounts for the possible slow absorption of information among market participants, which can make the instruments $\overline{MPI}_t^i$ autocorrelated. The four residuals $\{MPI_t^i\}_{i=1}^4$ constitute our final set of monetary policy instruments.

Figure 10 plots the four instruments both before and after adjustments for information and autocorrelation. The correlation between the unadjusted and the fully adjusted instruments is between 77 and 78 percent for all the four FRAs. The R^2 of the first-stage regression of the change in the money market rate on the four instruments is 10 percent.

Figure 11 shows the estimates from a local projection of the money market rate on an initial increase in the same rate, instrumented with our monetary policy instruments. We include one lag of the money market rate as well as month dummies and a second order polynominal in time.

C Additional figures

Figure 9: Interest rates.

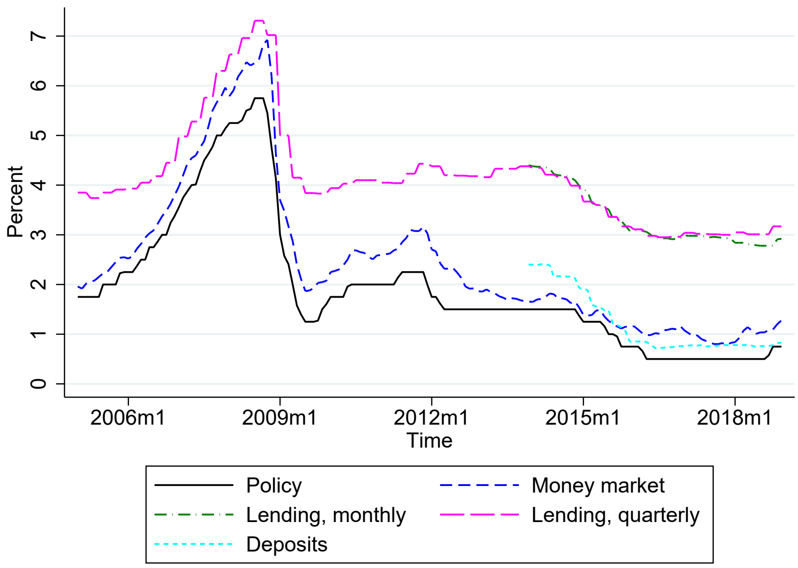

Notes. The figure shows interest rate series over our sample period 2006 − 2018. The policy rate is the interest rate on banks' overnight deposits in Norges Bank. The money market rate is the 3 month NIBOR rate. The monthly lending and deposit rates are the interest rates on all outstanding repayment loans and on deposits, respectively, for the household sector, both from a sample of banks and mortgage companies (these rates are only available starting in December 2013). The quarterly lending rate is from a full count of banks' and mortgage companies.

Figure 10: Monetary policy instruments.

(a) FRA1

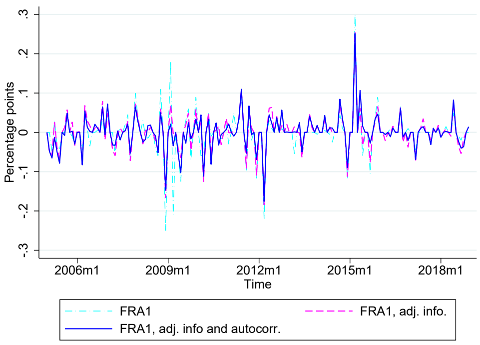

(b) FRA2

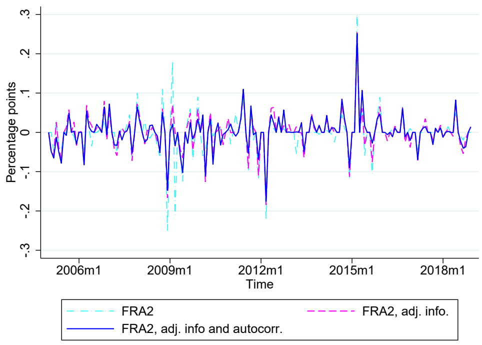

(c) FRA3

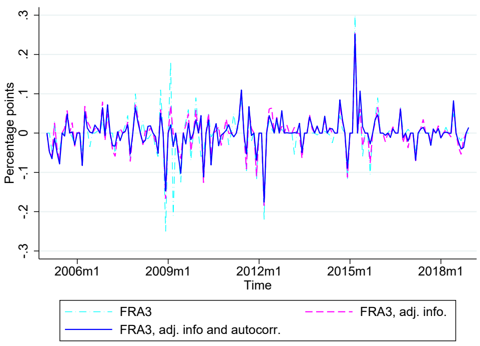

(d) FRA4

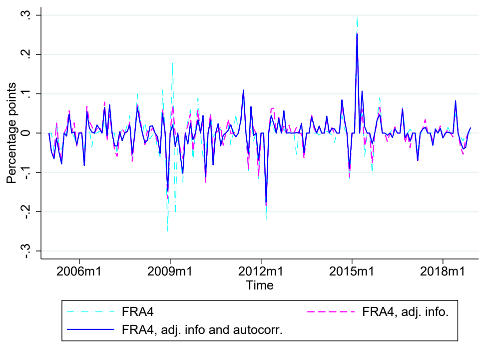

Notes. The figure shows the four monetary policy instruments used for the empirical specification in section 3, before and after adjustments for, respectively, information effects and autocorrelation.

Figure 11: Interest rate after initial change.

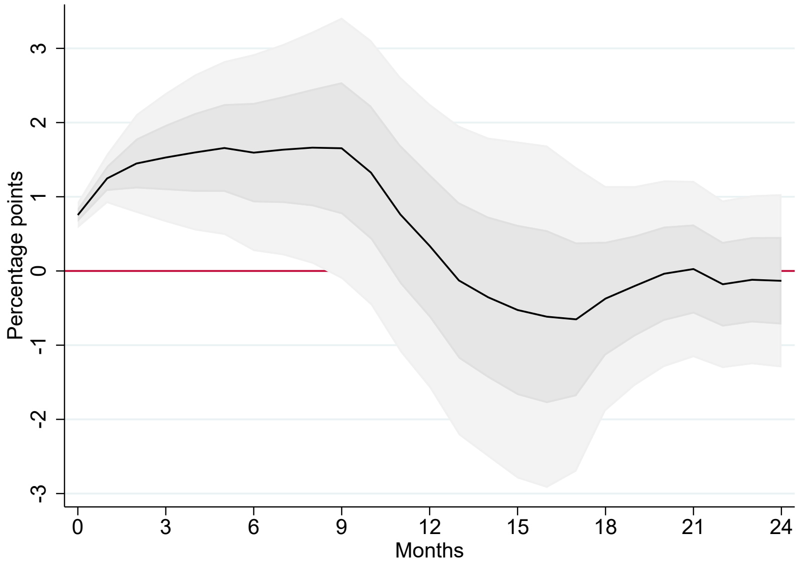

Notes. The figure shows the result of the local projection of the 3 month money market rate on the initial change $t = 0$ in the instrumented money market rate. The regression includes month dummies and a second order polynominal in time. See section B.

Figure 12: Pass-through of money market rate to lending and deposit rates.

(a) Lending rate

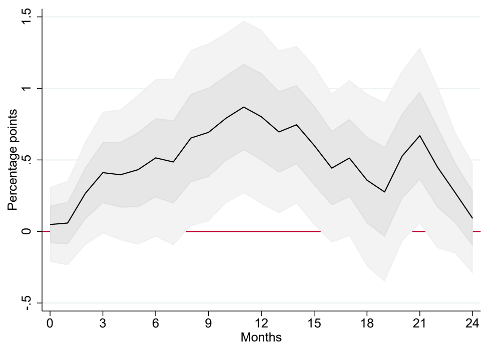

(b) Deposit rate

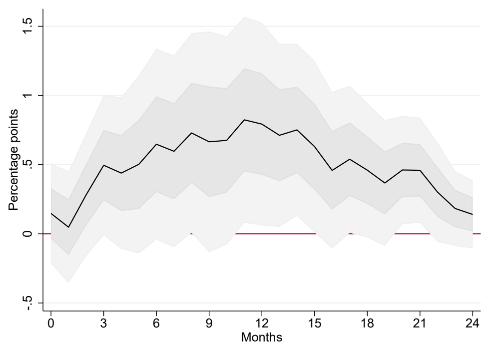

Notes. The figure shows the pass-through of an initial change in the 3 month money market rate to the lending and deposit rates, respectively, for the period December 2013 to December 2018. See section **??**.

Figure 13: Consumption response to 1 p.p. increase in interest rate, by interest exposure and month.

(a) Months 1

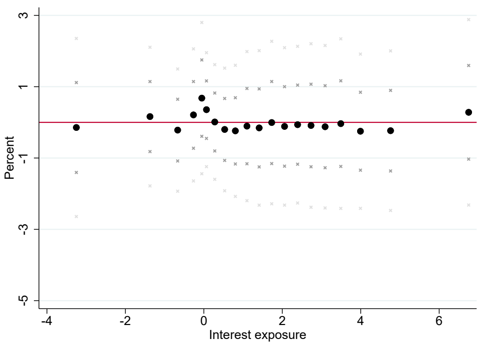

(b) Months 2

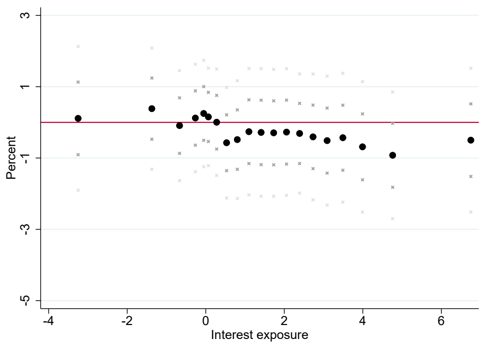

(c) Months 3

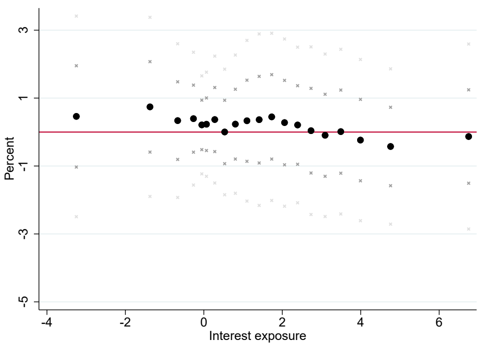

(d) Months 4

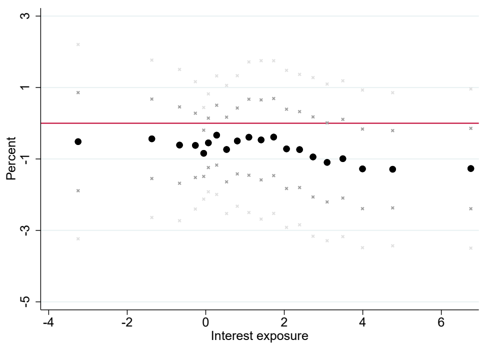

(e) Months 5

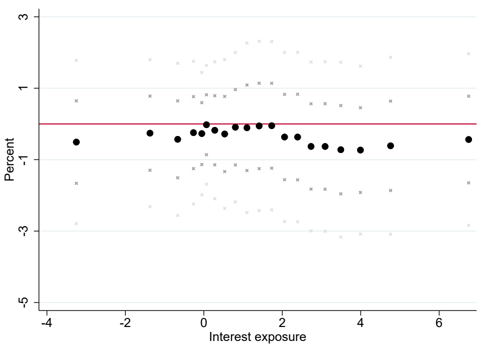

(f) Months 6

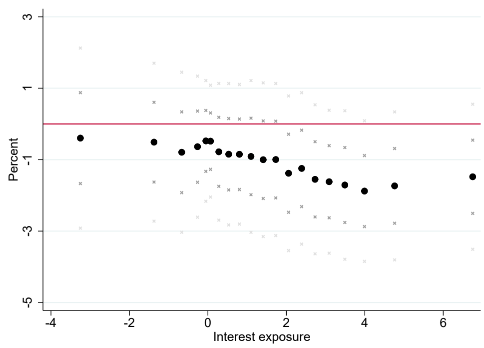

Notes. The figure shows the results from regression 2 by time horizon and quantiles (20 ventiles) of households based on interest exposure (debt subtracted bank deposits and divided by income). Each subfigure shows a separate horizon of the local projection. Month 0 is the month of the instrumented interest rate change. The black dots show the coefficient β_g^h of the regression for a particular ventile g and horizon h. The median exposure for each of the ventiles is shown on the horizontal axis. The vertical axis shows the estimated response of consumption, relative to its level in the three months before the change in the interest rate, in percent of income.

Figure 14: Consumption response to 1 p.p. increase in interest rate, by interest exposure and month.

(a) Months 7

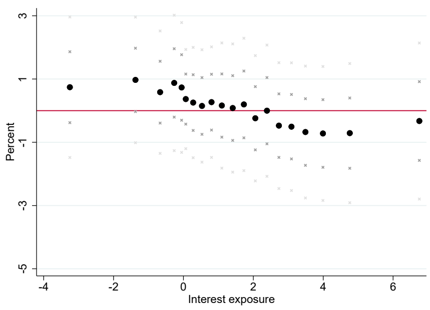

(b) Months 8

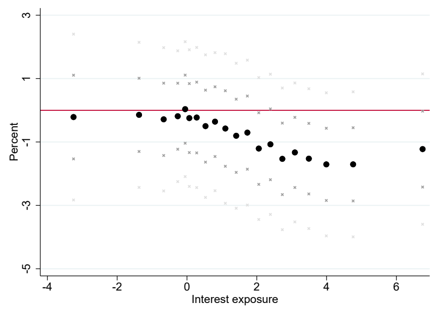

(c) Months 9

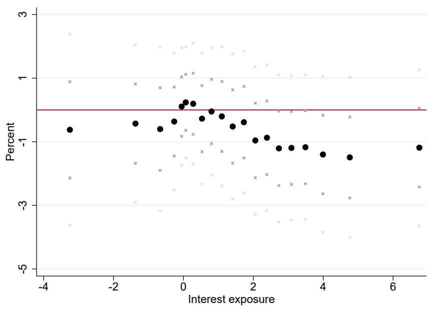

(d) Months 10

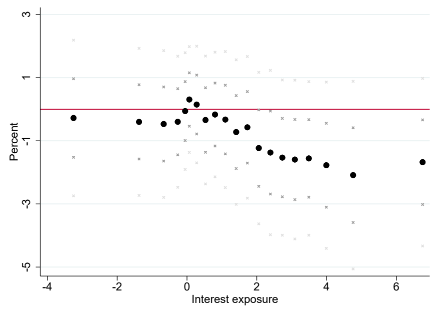

(e) Months 11

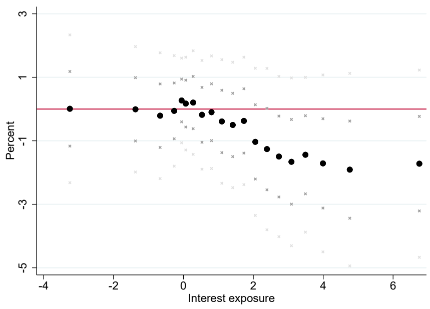

(f) Months 12

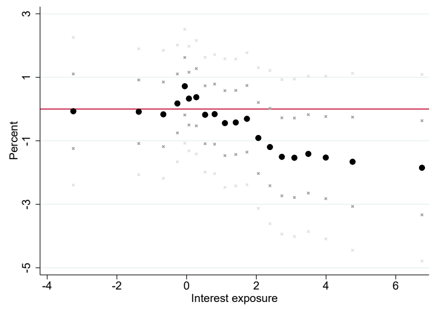

Notes. The figure shows the results from regression 2 by time horizon and quantiles (20 ventiles) of households based on interest exposure (debt subtracted bank deposits and divided by income). Each subfigure shows a separate horizon of the local projection. Month 0 is the month of the instrumented interest rate change. The black dots show the coefficient β_g^h of the regression for a particular ventile g and horizon h. The median exposure for each of the ventiles is shown on the horizontal axis. The vertical axis shows the estimated response of consumption, relative to its level in the three months before the change in the interest rate, in percent of income.

Figure 15: Consumption response to 1 p.p. increase in interest rate, by debt-to-income (DTI) and month.

(a) Months 1

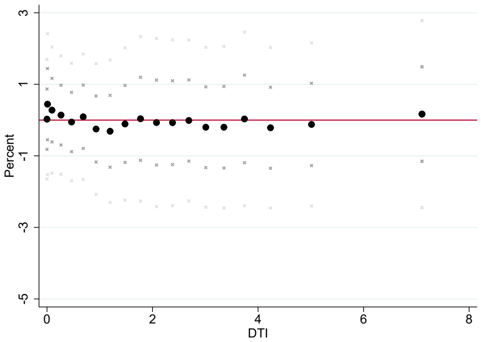

(b) Months 2

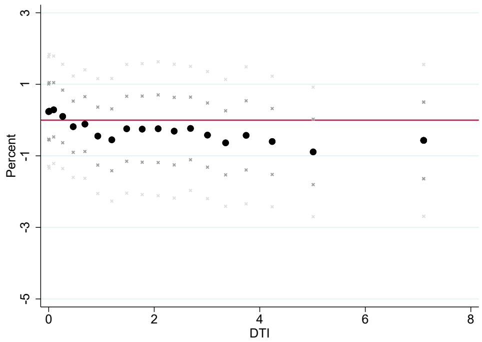

(c) Months 3

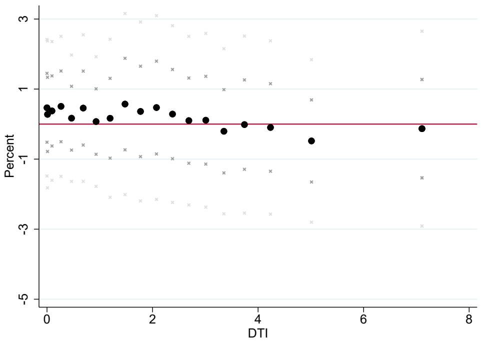

(d) Months 4

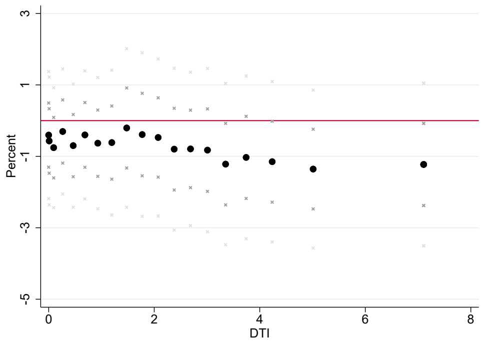

(e) Months 5

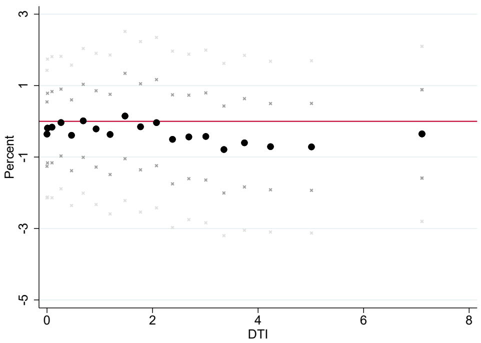

(f) Months 6

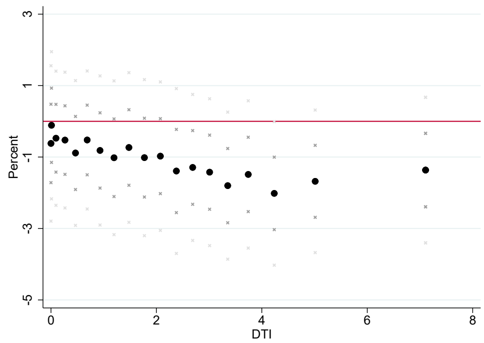

Notes. The figure shows the results from regression 2 by time horizon and quantiles (20 ventiles) of households based on debt-to-income (DTI). Each subfigure shows a separate horizon of the local projection. Month 0 is the month of the instrumented interest rate change. The black dots show the coefficient β_g^h of the regression for a particular ventile g and horizon h. The median exposure for each of the ventiles is shown on the horizontal axis. Because all the households in the bottom two ventiles have zero DTI, they are lumped together in the regression. The vertical axis shows the estimated response of consumption, relative to its level in the three months before the change in the interest rate, in percent of income.

Figure 16: Consumption response to 1 p.p. increase in interest rate, by month.

(a) Months 7

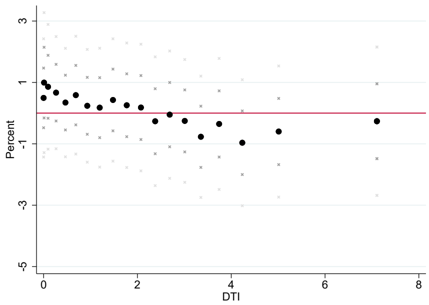

(b) Months 8

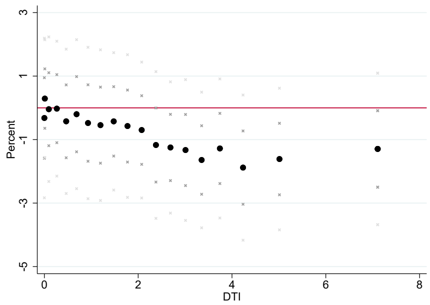

(c) Months 9

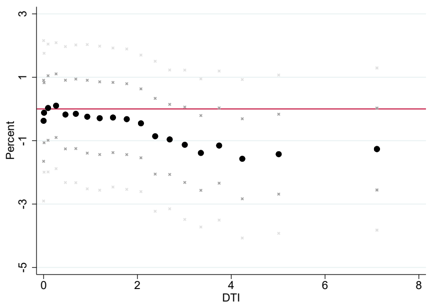

(d) Months 10

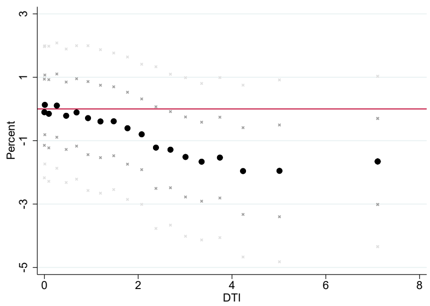

(e) Months 11

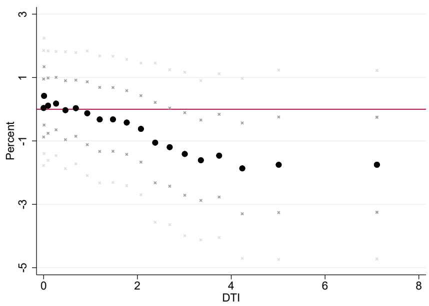

(f) Months 12

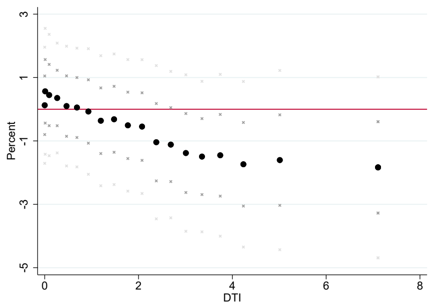

Notes. The figure shows the results from regression 2 by time horizon and quantiles (20 ventiles) of households based on debt-to-income (DTI). Each subfigure shows a separate horizon of the local projection. Month 0 is the month of the instrumented interest rate change. The black dots show the coefficient β_g^h of the regression for a particular ventile g and horizon h. The median exposure for each of the ventiles is shown on the horizontal axis. Because all the households in the bottom two ventiles have zero DTI, they are lumped together in the regression. The vertical axis shows the estimated response of consumption, relative to its level in the three months before the change in the interest rate, in percent of income.

Figure 17: Consumption response to 1 p.p. increase in interest rate, by interest exposure and month. Second year.

(a) Months 13

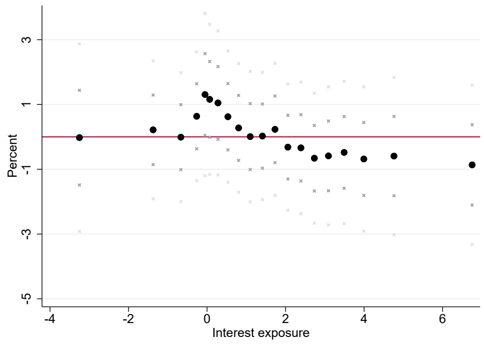

(b) Months 14

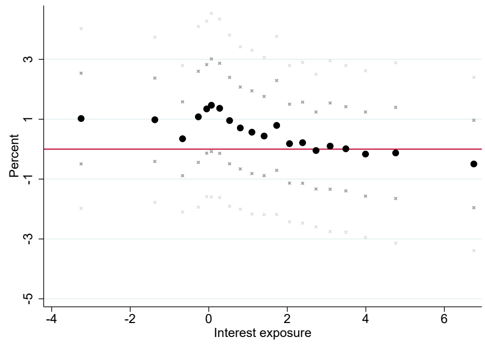

(c) Months 15

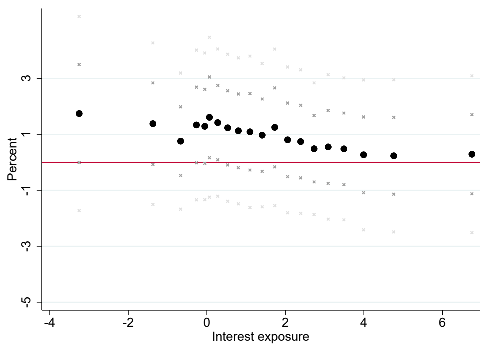

(d) Months 16

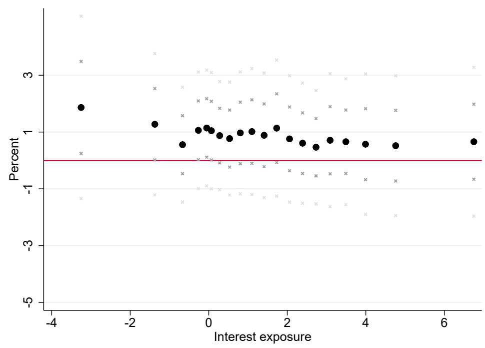

(e) Months 17

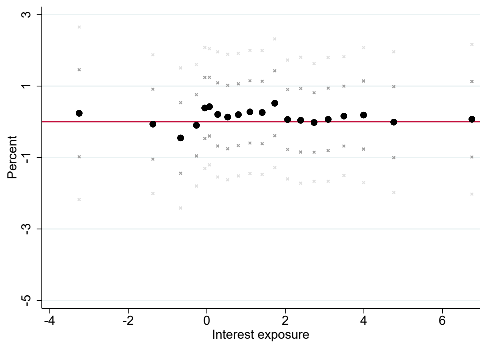

(f) Months 18

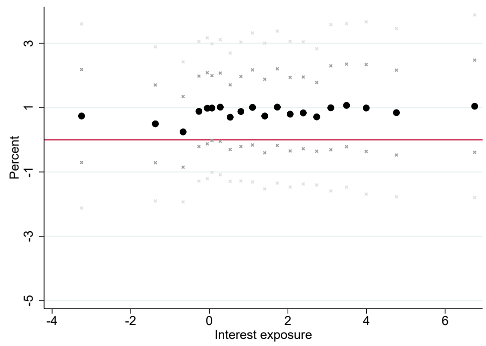

Notes. The figure shows the results from regression 2 by time horizon and quantiles (20 ventiles) of households based on interest exposure (debt subtracted bank deposits and divided by income). Each subfigure shows a separate horizon of the local projection. Month 0 is the month of the instrumented interest rate change. The black dots show the coefficient β_g^h of the regression for a particular ventile g and horizon h. The median exposure for each of the ventiles is shown on the horizontal axis. The vertical axis shows the estimated response of consumption, relative to its level in the three months before the change in the interest rate, in percent of income.

Figure 18: Consumption response to 1 p.p. increase in interest rate, by interest exposure and month. Second year.

(a) Months 19

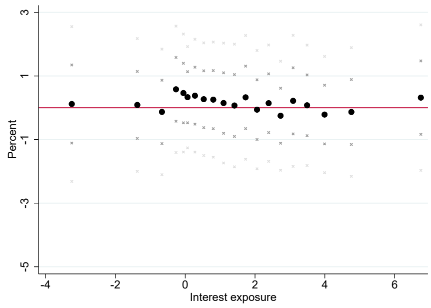

(b) Months 20

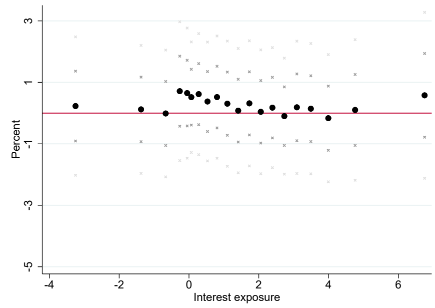

(c) Months 21

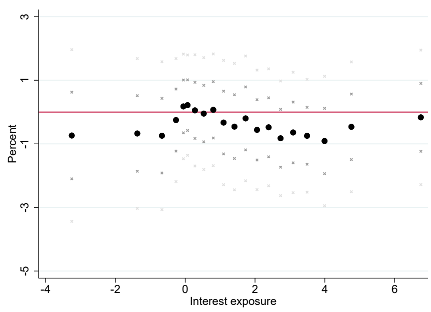

(d) Months 22

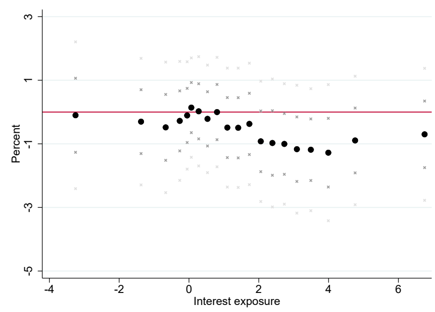

(e) Months 23

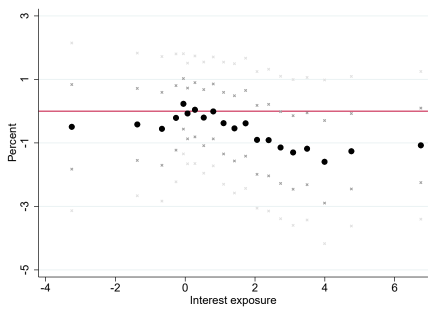

(f) Months 24

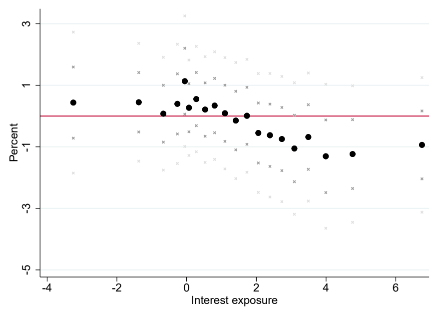

Notes. The figure shows the results from regression 2 by time horizon and quantiles (20 ventiles) of households based on interest exposure (debt subtracted bank deposits and divided by income). Each subfigure shows a separate horizon of the local projection. Month 0 is the month of the instrumented interest rate change. The black dots show the coefficient β_g^h of the regression for a particular ventile g and horizon h. The median exposure for each of the ventiles is shown on the horizontal axis. The vertical axis shows the estimated response of consumption, relative to its level in the three months before the change in the interest rate, in percent of income.

Figure 19: Consumption response to 1 p.p. increase in interest rate, by month and percentiles of interest exposure.