What Do 12 Billion Card Transactions Say About House Prices and Consumption?

Knut Are Aastveit* Jesper Böjeryd Magnus A. H. Gulbrandsen

Ragnar Juelsrud Kasper Roszbach

December 2024

PRELIMINARY - DO NOT DISTRIBUTE

Abstract

Changes in house prices can drive consumption, investment and the economy at large. Using close to the universe of debit card and electronic invoice payment data linked with relevant administrative background data for Norway, we estimate the dynamic response of consumption to a regional house price shock generated by the 2014-15 oil price decline. We isolate the effect of housing wealth by exploiting regional variation in exposure to the oil sector and studying government workers who face no change in unemployment or income risk and have stable and centrally determined wages. While a $1 drop in housing wealth causes an average decline in expenditures of $0.02 over a 24-month period, we find substantial heterogeneity across consumption components and population groups. Durable goods such as cars and furnishing respond strongest while semi-durables such as clothing are less affected and essential goods such as food and beverages do not react. Consistent with housing playing a role as collateral, the response of consumption increases with household leverage and decreases in liquidity and age. Our findings provide new evidence on the dynamics and the mechanisms through which house price shocks translate to changes in consumption.

JEL classification: D12, E21, E32, E44, E60.

Keywords: House prices, housing wealth, consumption, house price elasticity, marginal propensity to consume, payments data;

*Böjeryd is at UCLA, Stockholm University and Norges Bank; Aastveit, Gulbrandsen, Juelsrud, and Roszbach are at Norges Bank, firstname.lastname@norges-bank.no. 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. We would like to thank Yann Cerasi for excellent research assistance, and Adi Aladangady, Dominic Cucic, Daria Finocchiaro, Martin Flodén, Andreas Fuster, Refet Gurkaynak, Steinar Holden, Thomas Jansson, Niels Johannesen, Karin Kinnerud, Jonas Niemeyer, Søren Leth-Petersen, Pascal Paul, Federico Ravenna, Yannick Timmer, Neeltje van Horen, and Karl Walentin and other participants at the Swiss Winter Conference on Financial Intermediation, the conference on Monetary Policy and Financial Intermediation: Learning from Heterogeneity and Microdata at Collegio Carlo Alberto, BI Oslo, Sveriges Riksbank, CEBRA and ABN Amro for valuable comments.

1 Introduction

Improving our understanding of how house prices affect households' consumption and what factors drive this relationship is important for at least three reasons. First, the size of the consumption response is consequential for the aggregate economy's sensitivity to house prices (Guerrieri et al., 2020; Benmelech et al., 2022) and the housing market as a potential driver or amplifier of business cycle fluctuations. Second, housing is the foremost wealth component and collateral type used by households to obtain credit; housing heterogeneity among households is therefore a likely driver of variation in their ability to smooth consumption and response to shocks in the economy. Third, how households alter their consumption can have noticeable implications for spillovers effects from housing markets to other sectors and regions.

While recent studies, e.g. (Browning et al., 2013; Mian and Sufi, 2014; Aladangady, 2017; Paiella and Pistaferri, 2017; Aruoba et al., 2022; Graham and Makridis, 2023) have created a clearer picture of the quantitative effect of housing wealth on consumption, neither the mechanism through which this occurs nor the dynamics of consumption responses across heterogeneous households and consumption types are well understood and direct empirical evidence is relatively scant.

In a standard model of permanent income with infinitely-lived households, house prices have no effect on consumption because any wealth effect is counteracted by a reverse effect on living costs. In a life-cycle model, however, homeowners anticipating a future house sale will experience a wealth effect from changes in house prices that can dominate the impact of expected future living costs (Sinai and Souleles, 2005; Campbell and Cocco, 2007). In models where borrowing is collateralized, changing house prices will have a greater influence on consumption through households' borrowing capacity. Consumers with binding credit constraints have higher effective discount rates, making them more responsive to fluctuations in current wealth from changing home values compared to the negative effect of higher future living costs. Berger et al. (2018) show that such models predict large consumption

responses to permanent house price shocks. They estimate this effect to approximately equal the marginal propensity to consume out of temporary income times the value of housing, but argue it also depends on the level and distribution of debt, the size of the house price shock, and credit supply. The size effects in Berger et al. (2018) are in line with recent micro data evidence in Mian et al. (2013) and Aladangady (2017) who find average MPC's of housing wealth between 0.05 and 0.07 among homeowners and larger among credit constrained households.

Analyzing the impact of housing wealth on consumption *empirically* typically faces two challenges. The first is an econometric one, as housing is likely to be jointly determined with consumption, and underlying economic conditions may drive both simultaneously. This makes it difficult to draw strong inferences about housing wealth's role in the absence of (quasi-)experimental evidence. The second is the significant data requirement: an accurate understanding of how housing impacts on consumption requires that the econometrician can truly observe the consumption of households while controlling for potentially confounding factors such as wealth, debt and liquidity.

This paper attempts to take a step forward in identifying the role that housing wealth plays in driving consumption by exploiting a quasi-experimental setting and rich administrative data from Norway. Together, these enable us to overcome the two above empirical challenges. In particular, we use a unique dataset with nearly all debit card transactions, electronic invoice payments from households to firms, such as monthly credit card bill payments, direct remittances and [1] for the universe of Norwegian residents over the period 2009-2018. Because home ownership in Norway is around 80 percent and only 15 percent of payments are done in cash, the consumption response we derive is a reasonable approximation of the economy-wide response to the house prices change. The payments data we use capture consumption at the individual-week-zip code level through BankAxept, the Norwegian national payment system, for 26 different consumption COICOP categories.[2] Through

[1] While we have access to cash withdrawals, we exclude them from our measure of consumption

[2] The Classification of Individual Consumption by Purpose, abbreviated as COICOP, is a classification

an anonymized key, we match the transactions data with several administrative registers that contain information about income, wealth, employment, family situation, housing, education and other background characteristics.

The second part of our empirical strategy exploits differences in the geographical exposure to the oil sector and the regional variation in home prices that resulted from the 2014–2015 oil price shock. We focus on government workers and compare those who reside in the Greater Stavanger region - where the Norwegian petroleum industry is concentrated - with those living in the rest of Norway. Government workers in the Greater Stavanger region experienced an unexpected fall in house prices as a consequence of the local impact of the global drop in oil prices. Civil servants' wages, however, are stable and centrally negotiated at a *national* level. They therefore provide us with a treatment and control group with different exposures to the changes in local home prices but unchanged unemployment and income risk. By taking advantage of household-level housing transactions and tax data, we obtain household-specific changes in housing wealth. This quasi-natural experimental setting paired with comprehensive consumption data enables us to address the econometric identification challenges researchers face and estimate the *dynamic* marginal propensity to spend (MPX) out of housing wealth. The granularity of the consumption data paired with comprehensive household information permit us to do so across both various population groups as well as consumption categories.[3]

The main identifying assumption underlying our analysis is that the oil shock (and the location where households reside) only affect government workers' consumption through its impact on local home prices. We address potential threats to this assumption. First, we saturate our regressions with a large set of household controls. Second, after confirming that government workers in affected and unaffected regions follow similar consumption patterns

developed by the United Nations Statistics Division to classify and analyze individual consumption expenditures incurred by households, non-profit institutions serving households and general government according to their purpose (United Nations, 2018).

[3]Laibson et al. (2022) discuss the mapping from marginal propensities to consume (MPC) that are typically derived in macro models into marginal propensities for expenditure (MPX), which incorporate spending on a durable stock.

before the oil price collapse, we establish that government wages are unaffected by the oil price shock. Third, we address a potential threat to the exclusion restriction. If unobserved factors affect government workers in the oil region differently concurrent with the oil price collapse, for example because income expectations deteriorate, then this could jeopardize our identification strategy. We address this by analyzing two alternative samples: home-owning retirees and renters. Home-owning retirees are unlikely to be influenced by changes in local labor markets but may be affected by shocks in home prices, while renters are likely influenced by local labor market conditions but not by changes in home prices - if rents only adjust slowly to home prices. Reassuringly, retired home-owners display changes in consumption very close to government workers in our main regressions, while renters in the oil region exhibit no change in expenditures.

We thus study the effects of changes in housing wealth using an instrumental variables method that compares a treated group of government workers who live in the Greater Stavanger area where oil is the dominant industry and large oil price shocks affect house prices, with a control group of government workers from regions not highly dependent on oil.

We present three main findings. First, we document a marginal propensity to spend (MPX) out of housing wealth of approximately two percent over a twenty-four-month period. This number is in the lower range of what prior studies have found. Robustness checks on renters and retirees provide further support for the estimated effects. The high frequency of our payments data reveal that the cumulative MPX we commonly observe in studies of consumption masks a substantial non-linearity over time. As a consequence of the gradual materialization of a house price decline, the MPX is high at the start of the house price lapse when consumption responds but price effects still are small. As price effects accrue, the contemporary MPX initially declines before rising again as the cumulative house price effect grows, while the cumulative MPX monotonically rises.

Second, the comprehensiveness of the microdata enable us to uncover how home values influence spending across household segments. We establish that the MPX out of housing

wealth tends to increase with alternative measures of leverage and decrease with liquidity. As home values decline, borrowing constraints tighten for households near the collateral limit, hindering them from borrowing against their homes to finance consumption. Controlling for age, wealth, and liquidity we observe that households with a loan-to-value ratio close to a legal cap exhibit a marginal propensity to consume out of housing wealth of roughly 4 percent, double that of the average household. We also find that the MPX levels off for households far into retirement.

Third, we document the presence of substantial variation in the adjustment of spending across consumption categories. Our findings indicate that durable goods such as cars and furnishings exhibit the strongest negative response to falling home prices, while semi-durables like clothing are less affected and essential goods such as food and beverages do not react. The variation we uncover has likely implications for the general equilibrium effects of house price changes as, for example, localness of spillover effects can lead to local augmented of shocks as opposed to geographically smoothing.

This paper is probably closest to Aladangady (2017), Guren et al. (2021) and Andersen et al. (2024). Aladangady (2017) uses U.S. Consumer Expenditures Survey data and finds a quarterly MPC out of housing of 0.047. Andersen et al. (2024) exploit payments data from one large Danish bank and estimate an MPC out of small stock portfolio innovations equal to 0.044 over the first year and 0.16 over three years.[4] Our results align more closely with the estimates of Guren et al. (2021) who estimate the MPC to be about 0.03. Graham and Makridis (2023) find an MPC for non-durable spending of 0.8 to 0.9 cents per dollar. Using higher frequency housing price data, we find an MPC out of housing that is 0.02 at very short horizons, then declines and reaches about 0.02 again after two years. Qualitatively, our findings are consistent with the predictions in Berger et al. (2018). They show that models with incomplete markets and collateralized borrowing imply larger consumption responses to house price changes that depend on the level and distribution of debt and assets as well as

[4]Both papers exclude housing expenditures from consumption. Aladangady (2017) uses median home values.

the size of the house price shock. The MPC we find is also smaller than in Mian et al. (2013) and Kaplan et al. (2020). One possible explanation for this is that we find small responses for consumption categories that are typically local, while imported durables like cars react most. General equilibrium effects from spillovers, which are included in Aladangady's estimates are therefore likely smaller in our context. Another explanation may be that several of the earlier mentioned papers estimate consumption at the county or zip code level during the 2006-2009 US housing collapse, which was accompanied by a steep fall in the stock market. The latter likely contributed towards larger MPC estimates. Our estimates are derived from *household level* expenditure data. In addition, around the time of the 2014-2015 oil price decline, the Norwegian stock market did not fall substantially and in fact closely tracked the S&P 500. Finally, Kaplan et al. (2020) cite variation in expectations about the permanency of the 2006-2009 house price decline as a potential cause of variation in the estimated MPCs.

The principal contribution of this paper is to the empirical literature on the consumption effects of housing wealth shocks (Campbell and Cocco, 2007; Carroll et al., 2011; Mian et al., 2013; Browning et al., 2013; Aladangady, 2017; Andersen and Leth-Petersen, 2020; Aruoba et al., 2022; Benmelech et al., 2022; Böjeryd et al., 2023). While the above papers use annual survey data or imputed consumption, we complement these studies by also analyzing the dynamics of highly granular consumption. Specifically, our paper differs from the above studies in two important ways. First, we have access to register-based, high-frequency expenditure data from close to the universe of bank and households. Our payment data also capture transactions from accounts at all Norwegian banks and thereby increasing representativeness for households with multiple bank relationships. These features allow us to map the full dynamics of the consumption response at a quarterly frequency, providing new insights into the adjustment path of household spending. Second, we show how different types of consumption react, providing a starting point for an analysis of how house price shocks spread through the economy.

Our findings also add to studies of the consumption effects of other wealth and income

shocks (Berger and Vavra, 2015; Paiella and Pistaferri, 2017; Meyer et al., 2020; Fagereng et al., 2021; Imas et al., 2022; Andersen et al., 2024) by providing dynamic MPC estimates out of a housing wealth that can act as reference points for similar studies on other wealth and income shocks. Finally, we add to the macroeconomic literature on housing and consumption by offering new microdata evidence that can inform the calibration of macro models (Kaplan and Violante, 2014; Guerrieri et al., 2020; Kaplan et al., 2020; Baker et al., 2023).

Our findings have significant implications for policymakers. Although there is a prevalent policy concern that fluctuations in home prices may lead to substantial contractions in consumer spending, our results suggest a somewhat smaller overall impact of home prices on consumption, with muted negative spillover effects. Comparing our findings with those in Aladangady (2017) and Mian et al. (2013) suggests that the general equilibrium effects of broader, national, shocks and accompanying stock market declines may stand for about half of the measured response in consumption. Our findings also suggest that policies to reduce the level of household indebtedness may not result in a substantial change in the aggregate MPC, unless specifically aimed at households with high MPCs. The aggregate consumption response is a function of the distribution of indebtedness in the economy (Eggertsson and Krugman, 2012). Only highly indebted households, near the LTV cap, and those with very low levels of liquidity exhibit a stronger response. In our setting, this constitutes a relatively small group of households.

The remainder of this paper is organized as follows. In Section 2, we describe the data and variables. Section 3 outlines our research design, the empirical methodology, and the identification challenges. Section 4 presents the causal effects of home price changes on household consumption and discusses possible mechanisms behind our results. Section 5 concludes.

2 Data and Sample Construction

This section describes our data sources, the construction of our regression sample, the definition of central variable and provides summary statistics.

2.1 Data Sources

We base our analysis om four datasets. The first dataset contains expenditure data for the full adult Norwegian population at the individual x week x zip code level. The data have been provided by NETS, an international provider of, among other things, the technical infrastructure for in-store and online payments in Scandinavia, the Baltic region and Switzerland. NETS has obtained these data through its services to BankAxept, a company jointly owned by the largest Norwegian banks, that hosts an open system for physical debit card and electronic payments. Any bank business in Norway can join BankAxept and nearly all non-financial firms are connected to it. BankAxept processes debit card transactions, electronic bill payments, direct remittances and cash withdrawals for the universe of Norwegian residents and banks. It accounted for approximately 71 % of the volume of total card transactions over our sample period, all electronic invoice payments and direct remittances run through banks that are connected to BankAxept. To avoid identification of sensitive expenditure types, such as payments to health care providers, unions and political organizations, the expenditure data has been aggregated from single payments into 26 different consumption categories based on COICOP, the Classification of Individual Consumption by Purpose system. COICOP has been developed by the United Nations Statistics Division to classify and analyze individual consumption expenditures incurred by households, non-profit institutions serving households and general government according to their purpose (United Nations, 2018). NETS does not process or register online purchases that are not paid for by debit card or through electronic invoicing and processed through the Visa/EC/Amex systems. Payment of credit card balances is typically done through electronic invoicing, so our

data do capture the volume these transactions' amounts but we lack data on the decomposition of credit card spending. Credit card spending is less common in Norway than in, for example, the U.S. as the debit card and online invoicing data contain items like car purchases and expenditures on home durables. Typical items that will likely be underrepresented in our data are hotel reservations and car rentals abroad, i.e., when a credit card is commonly required by the recipient.

The second dataset contains administrative background data on all Norwegian individuals at an annual frequency and has been provided by Statistics Norway (SSB). We mainly use data from individual tax returns and register-based employment statistics. Since Norway levies both income and wealth taxes, tax filings data provided to SSB by the Norwegian Tax Administration provides a complete and comprehensive overview of the financial position of Norwegian households. It contains granular information about income sources, including information about wage income, business income, capital income and transfers, in addition to a breakdown of household balance sheets including bank deposits, the taxable value of primary and secondary housing as well as debt. It also includes *household* IDs, which we use for aggregation as described in Section 2.1. The employment register contains detailed information about employer-employee relationships and is administered by the Norwegian Labour and Welfare Administration and made available to SSB at an annual frequency.

Our third dataset contains detailed information about real estate holdings from the Norwegian Mapping Authority that was cleaned and structured by Ambita, a company that collects data on the real estate market. It contains real time data on all home ownership transactions and the selling price of homes and apartments. Finally, our fourth data source comes from Eiendom Norge and includes regional house price indices for a large number of Norwegian regions. We use the house price indices in combination with the the value of an individual's house to compute changes in the value of housing wealth at the individual level. We describe this procedure in more detail in Section 2.3. Figure 1 displays the evolution of the house price indices, where the red lines capture the different house price indices available

to us *within* the Greater Stavanger region.

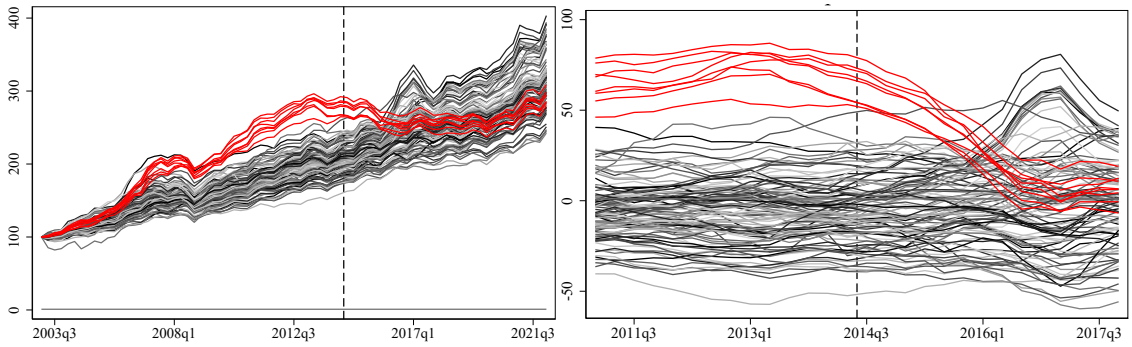


Figure 1: The left graph displays local home price indexes for different housing markets across Norway. The Greater Stavanger markets are marked in red and the rest of Norway in gray. The right graph shows the same indexes, detrended by the average year and quarter index. Source: Eiendom Norge (Real Estate Norway).

2.2 Sample Construction

For our analysis, we focus on households in which at least half of the adult members work in the government sector in both 2013 and 2014. This implies, for example, that families with a young adult will only be included in our group of "government worker households" if at least two members are civil servants. In Appendix A we describe in detail how we define the government sector based on statistical sector codes.

Our initial sample consists of 281,000 government worker households. To be able to observe a housing wealth effect, we condition inclusion in the benchmark regression sample on households being homeowners at the beginning of 2012. We additionally require that income from self-employment makes up less than 5% of a household's total income to avoid any noise from workers who rely on earnings that may be affected negatively by the oil price shock. After this, our final sample consists of about 206,000 households.

In our cross-sectional analysis, we compare outcomes in 2012 with outcomes post-oil price decline. We have 1,844,835 time x individual pairs over an 11 quarter period starting in 2014-Q3 and ending in 2017-Q1 for which we observe all card and invoicing transactions.

2.3 Variable Definitions

In this subsection, we describe and define the key variables we use in our analysis. Because we work with raw transaction-level data obtained from a payment provider, we need construct our our measure of household consumption. In doing so, we follow two principles. First, we apply COICOP, the global standard for classifying consumption into expenditure groups. Second, we inform our measure of consumption by definitions used in the literature, to facilitate a quantitative comparison of MPXs. However, because our expenditure data are more granular and available at higher frequency than in earlier work, we will study the dynamic response of consumption components that imply consumption in our analyses may differ from its definition in the national accounts or other papers. To provide support for our choices, we compare different components our expenditure based consumption with survey data Statistics Norway uses in its construction of consumption in the national accounts. All monetary values we present in this section are nominal Norwegian Krone (NOK) amounts.[5]

Demographic information For each household, we compute the average age of all tax-filing members and the number of individuals within the household who work in the government sector as defined in Appendix 1. We also collect the municipality, the number of children aged below 18 years, and several other demographic characteristics.

Consumption In the main analysis, we focus on four different, intuitively relevant notions of consumption: total consumption, food and beverages, furnishing, and vehicles. The latter three are quantitatively important subcategories of consumption as well as salient examples of non-durable, semi-durable, and durable consumption. To ensure anonymity of the subjects, consumption expenditures have been aggregated into 26 COICOP categories before being delivered to us. Using thee 26 categories, we define total consumption as the sum of all expenditure categories less payments to the public sector. Payments to the public sector are, together with other smaller, non-classifiable expenditures into a residual category

[5]In 2015, the USD/NOK exchange rate fluctuated between 7.3 and 8.8.

"Consumption, miscellaneous" that we do not use in our analyses.

When we compare our data to consumer surveys conducted by Statistics Norway (SSB), Table 1 shows our measure of total expenditures is somewhat smaller than that obtained in the SSB survey, but within a 1.8–15% margin depending on the region we consider. Furnishing is similar (1.8–10%). Our measure of food and beverages is about 33–46% greater and vehicles 27–36% less (see Table 1). These differences are likely to be caused by differences in the definition of consumption categories between the SSB survey and our expenditure data. The SSB survey questions include, for example, purchases of used cars while our measure will exclude these when a car is transacted between private individuals and not invoiced or paid by debit card.

Table 1: Comparison to consumer surveys

	Total consumption	Vehicles	Furnishing	Food & beverages
Microdata	397776	15428	24936	91880
All Norway	435507	43226	24495	63146
Rural	404862	57742	24224	65185
Urban	445068	42282	25962	64175
Agder & Rogaland	440794	44203	26950	64417
Two-person households	466821	50358	27694	69030

Note: This table compares the expenditures in the microdata to survey data collected by Statistics Norway and used for the national account. We use the 20123 wave and split up consumption by different characteristics. The survey expenditure categories do not coincide exactly with our definitions. In the survey, "vehicles" refer to consumption group 071 ("vehicle purchases") which also includes used vehicles, "furnishing" refers to consumption group 05 ("furniture, household items, and maintenance"), and "food and beverages" is the sum of groups 01 and 02 ("groceries and alcohol free beverages" and "Alcohol beverages and tobacco"). Urban is defined a non-rural but excluding Oslo, Bergen and Trondheim.)

Household income and balance sheet variables We measure the value of a household's house by the implied market value of the household's primary house at a given point in time using the estimated market value reported by the tax authorities in 2012 and the evolution of

the local house price indices post 2012. The loan to value (LTV) of a household is obtained as total household debt over the sum of the estimated market values of primary and secondary housing. We calculate the debt to income (DTI) ratio by dividing total household debt with after-tax income. Finally, we compute liquid wealth to income (LTI) by dividing bank deposits by after-tax income.

2.4 Summary statistics

Table 2 presents descriptive statistics for our main regression sample. Average total consumption equalled approximately 99,000 NOK per household in 2012q4, with a median of about 86,000 NOK. Out of the three subcategories of consumption, food and beverages are most important with average expenditures per quarter of approximately 23,000 NOK. Average expenses on furnishing and vehicle purchases are approximately 6,000 and 4,000 respectively. The average Norwegian household owns a primary house with a value of 2,7 million NOK (mNOK), with a standard deviation of 1,3 mNOK. The LTV distribution is very skewed, with a mean of 0.90, a median of 0.48 and a standard deviation of 220. In our main sample the average (adult) age in a household is 47 years with 1.42 members working in the government sector. Approximately five percent of our sample resides in the Greater Stavanger region.

3 Research Design

This section presents our research design. We discuss the empirical setting and underlying assumptions, and provide a series of robustness tests to address potential challenges to our identification strategy.

Table 2: Summary Statistics

	Mean	Median	Std. dev	N
Consumption (in 2012q4) in NOK				
Food and beverages	22 970	20 834	14 805	205 853
Furnishing	6 234	3 429	10338	205 853
Vehicles	3 857	0.00	24883	205 853
Total consumption	99 444	86 601	66 853	205 853
Household income and balance sheet variables				
House value (primary)	2 704 311	2 454 016	1 308 990	205 853
Loan to value (LTV)	0.90	0.48	167.69	205 853
Debt to income (DTI)	2.48	2.20	29.08	205 852
Liquid wealth to income (LTI)	6.96	3.12	13.96	205 852
Demographics				
Average age in the household	46.56	45.00	11.10	205 853
No. individuals in gov.sector in household	1.42	1.00	0.63	205 853
Greater Stavanger dummy	0.05	0.00	0.22	205 853

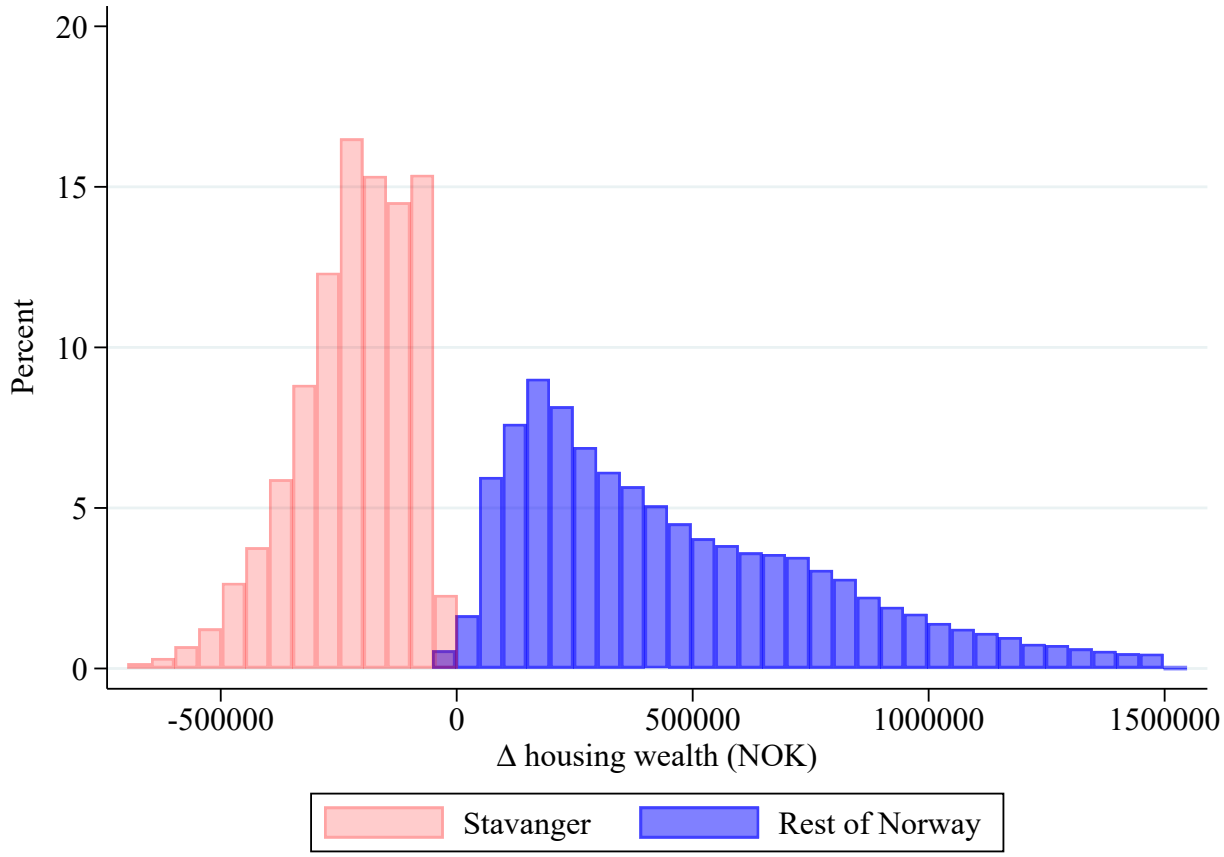
Notes: This table presents summary statistics for the data used in the benchmark analyses and regressions. Food and beverages, furnishing and vehicles refers to the NOK amount spent per quarter on those goods. Total consumption is the total amount spent on all consumption categories per quarter, in NOK. The loan to value ratio (LTV) is defined as a household's total debt divided by the value of its primary housing. The Debt to Income ratio (DTI) is total debt to total income for a household. Liquid Wealth to Income (LTI) is defined as total bank deposits on December 31 and divided by (total gross annual income/12).

3.1 Experimental Setting

In the ideal experimental setting to estimate the dynamic effect of house price changes on consumption, the econometrician would observe one-off, random, unpredictable variations in housing wealth at the household level. The quasi-experimental setting we exploit approximates the ideal conditions because it provides us with household level variation in housing wealth induced by the sharp decline in the oil price in 2014. As documented in Section \2.1, the economic downturn in the petroleum sector and regions dependent on the petroleum sector that followed the 2014 price oil fall led to substantial drops in local home prices. Figure 1 exhibits the evolution of home price indices in Greater Stavanger—the most exposed region to the oil price collapse—and other areas of the country. We combine these local house price changes with the geographical distribution of households over different local housing markets. Doing so links population densities to each of the curves in Figure 1. Figure 2 depicts the implied distribution of house prices *changes* at the *household* level over the period 2015-2016. We employ this variation in house prices in our research design. In Section 4.1 we describe how we use an IV approach to account for the fact that exposure to house price shocks may not be exogenous.

In our empirical strategy we exploit differences in the geographical exposure to the oil sector and the regional variation in home prices that resulted from the 2014–2015 oil price shock. We focus our analysis exclusively on government sector workers such that our analysis compares civil servants who live in the region around Stavanger, where the petroleum industry is the most important sector of the economy, with civil servants elsewhere in Norway. Government workers in the Greater Stavanger region experienced an unexpected fall in house prices as a consequence of the local impact of the global drop in oil prices. Civil servants' wages, however, are stable because they are centrally negotiated at a *national* level through collective bargaining between trade unions and employer federations. Additionally, government jobs in Norway are highly secure with de facto zero unemployment risk. Although a significant portion of national fiscal spending is funded by a special tax on petroleum

Figure 2: Distribution of change in home prices



Note: This figure shows the distribution of home price declines over the period 2015-2016 at the household level. Red bars depict home price growth for households in the oil intensive region, grey bars show home price growth for households living outside the oil region.

activities, the distribution of these funds is governed by a fiscal policy rule and set at 3% of the return on the portfolio of the Government Pension Fund Global (the "oil fund") and thus only negligibly influenced by the current *inflow* into the fund. Finally, government spending is not systematically related to a location's exposure to the petroleum industry as a result of which any changes in either current or prospective spending by the government on civil servant salaries will apply uniformly across the country. Unlike private sector workers, civil servants are thus exposed to fluctuations in house prices through their home ownership

without facing any simultaneous changes in income prospects. These circumstances provide us with a natural group of treated - and control workers with varying exposure to house price changes and unchanged unemployment and income risk.

3.2 Verifying our assumptions

As a first test of of our conjecture that government workers in this oil intensive region suffered losses to the value of their homes and spent less relative to the rest of Norway following the 2014 oil price shock. For this purpose, we run a standard difference-in-difference regression for government workers' spending and home values using the oil region dummy as a treatment indicator, i.e., we estimate:

$$y_{i,t} = \beta \, \text{post}_t \times \text{oil region}_i + \eta \, X_{i,t} + \varepsilon_{i,t}, \tag{1}$$

where $y_{i,t}$ is the logarithm of either household i's home value, total consumption, or labor income in year t. The vector of controls $X_{i,t}$ contains household mean age and size, and fixed effects for households, housing markets, and years. In these regressions, we ignore any effects that stem from individual (housing) asset holdings and focus on pure locality effects. Additionally, to verify a central assumption underlying our empirical strategy, we also check if government workers in the Greater Stavanger area indeed experienced no change in labor income. We therefore run a similar d-i-d regression for total labor income. Table 3 shows government workers in the oil intensive region suffer a significant drop in relative home values of 18 percentage points. They also cut back on consumption albeit the change is only marginally significant, reflection large variation in the consumption effect. Importantly, labor income is not changing differently across regions. The estimated coefficient in column (3) is -0.0060, with a standard error of 0.016. This soundly confirms our assumption that government workers are on the same earnings path in the Greater Stavanger region and the remainder of Norway.

Table 3: Difference-in-difference model results

	(1) Log(House Val)	(2) CapInc/FinAssets	(3) Log(Labor Inc)	(4) Log(SE Inc)
Post x Oil region	-0.184***	0.000150	-0.00604	-0.00194
	(0.0398)	(0.000202)	(0.0156)	(0.00973)
R^2	0.916	0.290	0.518	0.533
Household FE	Yes	Yes	Yes	Yes
Region FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Nobs	1,472,945	1,272,366	1,472,945	1,472,945

Note: This table presents the estimates of (1). SE Inc is Income from self employment Errors are three-way cluster-robust at the level of household ID, housing market, and year. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.3 Threats to the research design

The key identifying assumption underlying our research design is that the oil price shock only affects government worker expenditures through its effect on home prices and not through other location-specific channels. Arguably, there are two potential threats that could lead to a violation of this assumption. First, there could be structural differences between our treated households and those in the control group, Second, while there are no differences between the treated and control households prior to the oil prices shock, there could be confounding shocks that differentially affect the treatment and control group. We investigate both issues below.

Structural differences between treatment and control

A potential concern could be that government workers differ fundamentally across regions in terms of their consumption patterns. For instance, different regions experience distinct business cycles, government workers residing in areas that are, relatively, booming could be on a higher consumption growth than workers residing in other areas. This could be due,

for example, to earnings being correlated with the local business cycle, or to peer effect (government workers in Stavanger consuming more because their non-government-worker neighbors are consuming more).

To examine this, we start our empirical analysis by estimating a flexible differences-in-differences model where we compare consumption in Greater Stavanger to the rest of the country both before and after the oil price collapse. To the extent that consumption evolves similarly prior to the oil price collapse, this would suggest that government workers in different areas have relatively comparable consumption patterns prior to the oil price collapse.

Because of the high frequency nature of our data, we face substantial within-household seasonality in consumption. To deal with high frequency volatility that would introduce noise into an evaluation of consumption responses, we remove the seasonality by computing the household's moving average consumption over the current quarter and the three leading quarters.

$$\widehat{c_{i,t}} = \frac{c_{i,t} + c_{i,t+1} + c_{i,t+2} + c_{i,t+3}}{4} \tag{2}$$

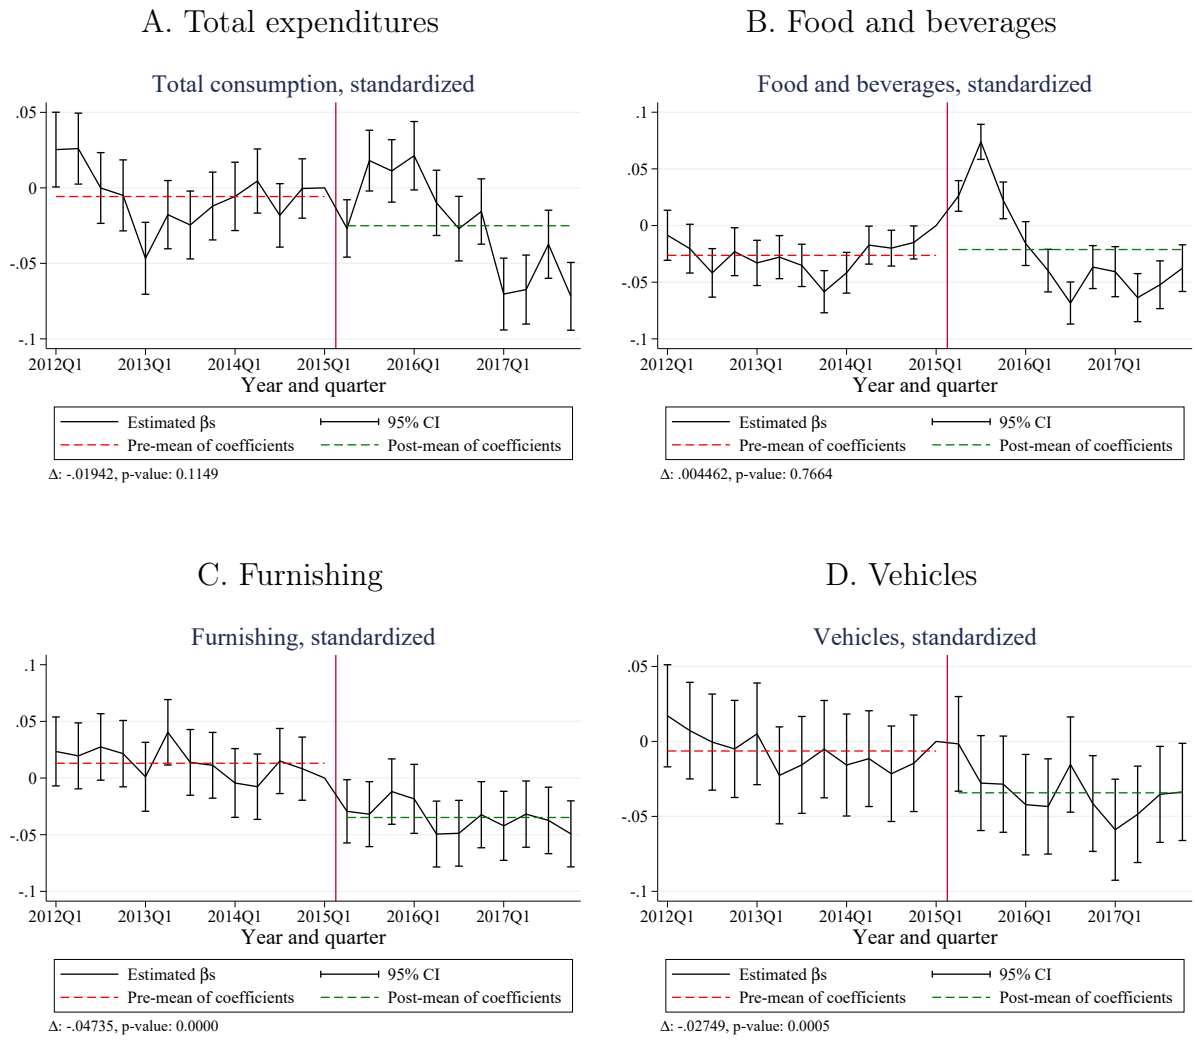
We regress this on a series of year-quarter dummies interacted with a dummy if the household lives in Greater Stavanger in 2014, excluding 2014Q4,

$$\widehat{c_{i,t}} = \sum_{\substack{\tau \in \{s,\ldots,S\}, \\ \tau \neq S_0}} \beta_\tau \, \mathbb{1}(t = \tau) \times oil\ region_{i,2014} + \eta_i + \eta_t + \varepsilon_{it}. \tag{3}$$

where $s = 2012Q4$, $S = 2018Q4$ and $S_0 = 2014Q4$. The regression also includes year-quarter and household fixed effects. Errors are two-way cluster-robust at the level of housing market-region.

Figure 3 displays a plot of the difference between consumption between the Stavanger area and the rest of Norway. The four panels make clear there is no persistent difference in consumption between the regions in the pre-period. Post 2015, however, spending on furnishing and vehicles (Panels C and D) fall significantly in Stavanger region.

Figure 3: Variation in consumption across time, location, and consumption category



Note: This figure presents the dynamic response of expenditures for different categories of consumer goods estimated using (3). The difference Δ and the following p-value under each legend refer to the difference between the red and green dashed lines, indicating the means of the coefficients in the pre- and post-periods, respectively. This is estimated using a difference-in-difference model similar to (3) but with the year-quarter dummies replaced by one post dummy.

Confounding shocks

Next, we investigate the presence of any confounding shocks that could differentially affect government workers after the oil price collapse. An example of such a shock could be that income expectations deteriorate for government workers in the oil region relative to government workers otherwise. We follow two approaches. First, we condition a model of earnings on a wide set of predetermined observables to ensure that the comparison is based on government workers that are similar based on observables. Second, we perform an additional analysis on an alternative sample of only home-owning retirees. While working adults' expected lifetime earnings can be affected by expectations of a future deterioration in labor market prospects, retirees lifetime incomes are governed by their pension entitlements and unaffected by the evolution of the local labor markets. At the same time, retirees are exposed to the same housing wealth shocks as working adults. Should we find that retirees display similar consumption responses as working adults, then this would assuage any concerns about confounding events.

4 Results

In this section, we present our main empirical analysis where we estimate and discuss the marginal propensities for expenditures of selected consumption goods. We also perform robustness tests discussed earlier and heterogeneity analysis to separate between the potential channels at work.

4.1 IV model

To obtain the MPXs out of housing wealth, we want to estimate

$$\Delta c^k_{i,t+\Delta t} = \beta^k \, \Delta hw_{i,t} + \eta \, X_{i,t} + \varepsilon^k_{i,t},$$

where β^k is the marginal propensity for expenditures of good k in period $t + \Delta t$ of a change in housing wealth in period t. In the absence of exogenous variation, one cannot estimate this model because of endogeneity between $c^k_{i,t+\Delta t}$ and $hw_{i,t}$, typically attributed to omitted variables such as income - and wealth expectations (Attanasio et al., 2009)). We avoid this by estimating an exogenous change in $hw_{i,t}$ for government workers in different housing markets in Greater Stavanger using an IV (2SLS) approach.

For each household, we observe the value of their primary home at the beginning of 2014 which we denote by $hw_{i,2014}$. Using quarterly home price indices from Eiendom Norge, we compute an estimated value of the home in each quarter over 2012–2018. Based on this baseline housing wealth, we then start in the fourth quarter of 2014 and compute each household's mean primary home value in that window, and denote it by $hw_{i,2014+\Delta t}$. We compute the housing wealth change within household i since 2014 and regress this difference on a dummy that captures if the household lives in the oil-intensive region ℓ (equation (4)) while controlling for household variables that could be correlated with the local change in home prices and drive the change in consumption.

$$\Delta hw_{i,2014+\Delta t} \equiv \frac{hw_{i,2014+\Delta t} - hw_{i,2014}}{hw_{i,2014}} = \delta \times oil\ region_{i,2014} + \quad \eta_1 X_{i,2014} + \alpha_1 + \varepsilon_{1\,i,\Delta t} \quad (4)$$

Here, *oil region* refers all households located in Greater Stavanger region as defined by Eiendom Norge. The vector of controls $X_{i,2014}$ contains labor income, capital income, gross wealth, liquid wealth, and a household's total debt in 2014, each with polynomials up to degree four. All control variables are standardized. The intercept α_1 is the average change in home prices across the country over the considered period excluding the Greater Stavanger market and is not explained by changes in the characteristics of local government workers.

The predicted change in housing wealth that we can estimate with (4) is thus the change in the value of households primary's housing due to the location's exposure to the oil price

fall. For households living outside the Greater Stavanger region, the *exogenous* change in housing wealth is zero by construction.

In the second stage, we estimate for each consumption good category k a linear model:

$$\frac{c_{i,2014+\Delta t} - c_{i,2014}}{hw_{i,2014}} = \beta_{\Delta t} \widehat{\Delta hw}_{i,2014+\Delta t} + \quad \eta_2 X_{i,2014} + \alpha_2 + \varepsilon_{2\,i,\Delta t} \tag{5}$$

where, the controls in $X_{i,2014}$ are the same as in the first stage. Our central parameter of interest $\beta_{\Delta t}$ should then be interpreted as "the expenditures of particular goods or services category increase by $\beta_{\Delta t}$ between periods 2014 and $2014 + \Delta t$ for every NOK rise in the household-specific housing wealth between 2014 and $2014 + \Delta t$". We multiply the outcome variables by four in order to express the coefficient in per-year terms instead of quarterly effects.

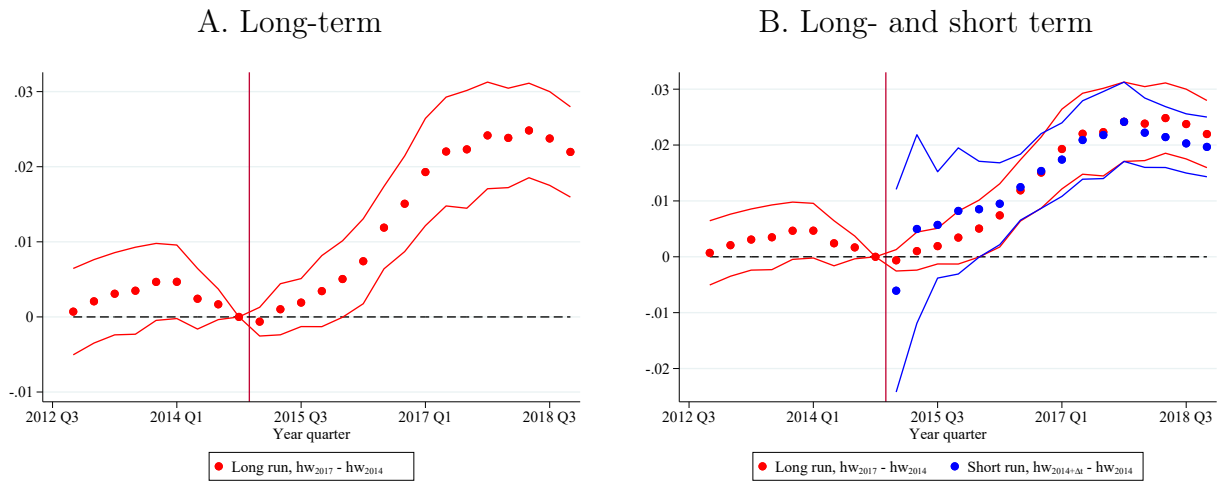
In the estimation of the baseline regression, we keep the benchmark year fixed at 2014 while moving the end-year forward by a quarter at a time. We start in 2014Q4–2015Q3, then continue with 2015Q1–2015Q4, etc until we reach the end of the sample. Here, the change in housing wealth will vary quarter by quarter and gradually build up. In a second regression, we set $\widehat{\Delta hw}_{i,2014+\Delta t} = \widehat{\Delta hw}_{i,2017Q3}$ for all Δt so as to obtain a "cumulative" MPX. The latter facilitates a comparison with earlier estimates of the MPX that typically rely on "final" house price shocks

4.1.1 The average consumption responses

In Figure 4 we show the MPXs of total expenditures and three categories of consumer goods: Food and beverages, Furnishing, and Automobiles. Each MPX should be interpreted as the annual expenditure effect. We do not observe homeowners' expectations about *future* home prices. Our main specification in equation (5), therefore implicitly assumed that any change in spending is due to the *current* reduction in home prices in Stavanger relative to the rest of Norway. Workers may, however, internalize that prices will diverge even further. To study the potential consequences of this, we provide a second set of estimates where we hold the

change in housing wealth fixed at the longest possible window we can consider with our data,

Figure 4: The dynamic marginal propensity to spend out of housing wealth shock

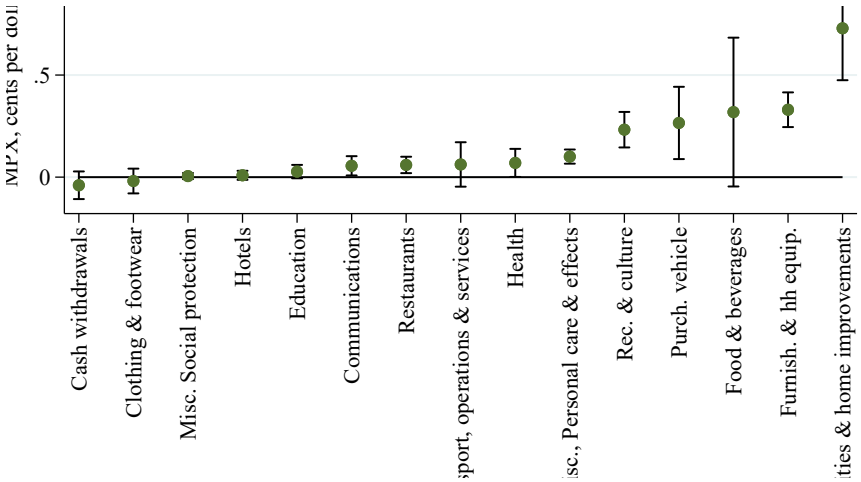


Note: This figure presents the dynamic response of total consumption expenditures for home-owning government workers. Each period refers to a 4-quarter episode that starts in the time on the horizontal axis. Short-term MPXs (blue) are estimated using changing home prices, while long-term MPX's (red) display the corresponding estimates when the household-specific housing wealth change is computed over the longest range. The confidence intervals are 90% and are cluster-robust at the housing market level.

i.e., the accumulated change over the period 2014–2018. These are displayed in red in Figure 4, panel A, while the original set of estimates is represented in blue. Total expenditures fall with decreasing home prices, as expected. The magnitude of this change is within the range of earlier estimates but in the lower range of them. Quantitatively, our estimates are closest to Disney et al. (2010), Guren et al. (2021), and (Böjeryd et al., 2023).

Early on in the process of falling house prices, consumption responses are about half the size when we use the longer-horizon housing wealth change compared to the current-housing-wealth changes (red versus blue), but the effects are identified more accurately. As we look further ahead, this difference gradually disappears by construction since the instant and accumulated changes in housing wealth eventually converge (panel B).

Figure 5: The cumulative MPX Across Consumption Categories



Note: This figure presents the total, i.e., the accumulated response of consumption for different expenditure categories. Each MPX refers to the total MPX over the period 2014q3-2018q3. The confidence intervals are 90% and are cluster-robust at the housing market level.

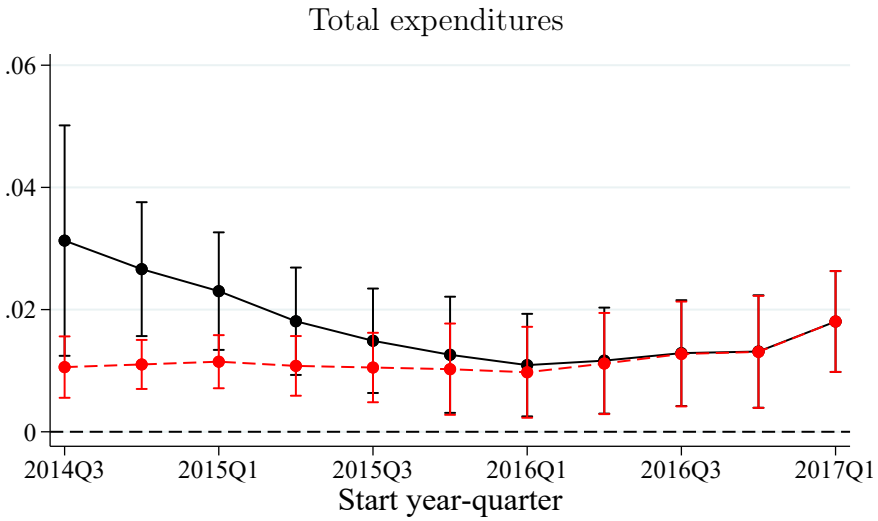
In Figure 5, we see that expenditures on Food and beverages show no sign of significant adjustment. In sharp contrast, housing wealth has an instant, large and significant effect on large expenditure durables. A one-NOK fall in housing wealth causes an reduction in (mostly new) car purchases of 0.08 NOK reduction a the beginning of the house price decline (not shown here) that gradually levels off to about 0.04 as the house price decline builds up over time. Simultaneously, we see that households sharply reduce expenditures on goods and services directly associated with housing, such as on furnishing (0.04) and home improvement (0.08). Unlike car purchases, the initial effect is smaller and builds up gradually over the course of time following the first house price decline (not shown here).

4.1.2 Robustness analysis

We perform two robustness checks. The key assumption underlying our analysis was that Stavanger government workers' income prospects do not change relative to other government

workers in Norway, but the former group experiences a big unexpected change in housing wealth. If income prospects do change due to the oil price shock, this would bias our results. Another group that is more isolated to income risk is retired homeowners. We define people as retired based on their labor earnings and social benefit payments linked to retirement.

Figure 6: The dynamic marginal propensity for spending out of housing wealth shock for retired homeowners (NOT UPDATED)



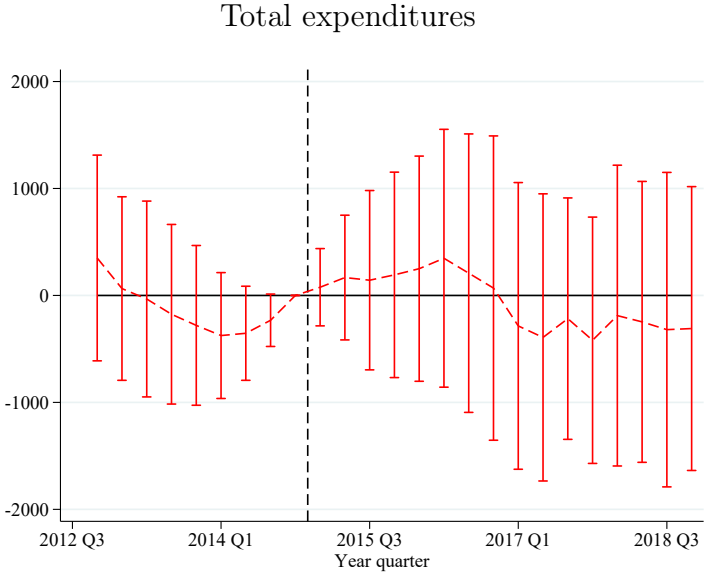
Note: This figure presents the dynamic response of total consumption expenditures for homeowning retirees. Each period refers to a 4-quarter episode that starts in the time on the horizontal axis. The black solid lines display the MPXs estimated using changing home prices, while the red dashed lines display the corresponding estimates when the household-specific housing wealth change is computed over the longest range. The confidence intervals are 90% and are cluster-robust at the housing market level.

We then rerun our main specification (5) on a sample of retired homeowners and contrast similarly between those living in Greater Stavanger and the rest of Norway. The results are presented in Figure 6.

The responses of retirees are also significant and in magnitudes greater than for the main sample of government workers. However, the estimates are not precise enough to tell if the results are significantly different.

In a second robustness we test if renters who work in government respond similarly or not. By the assumptions we have made, government workers who rent should be isolated from both income and housing wealth risk. Because renters have no real estate assets, we cannot compute a standard MPX for them since they have no change in housing wealth. We therefore re-estimate our benchmark model 1 using the total expenditures of renters and present the results in Figure 7.

Figure 7: Consumption response for renters



Note: This figure presents the dynamic response of renters' total expenditures (MPX) as estimated by equation 5. Each period refers to a 4-quarter episode that starts in the time on the horizontal axis. The confidence intervals are 90% and are cluster-robust at the housing market level.

Figure 7 shows that renters do not experience any significant downward shift in total expenditures. This provides support for our assumptions as renters would be exposed to other, confounding, shocks while being insulated from changes in house prices. The absence of any consumption response by renters thus provides strong support for the notion that there are no confounding shocks that drive consumption downwards for government workers.[6]

[6]Our test on renters has less power, however, because the renters sample is substantially smaller than

4.2 Why do home prices affect consumption?

Our goal in this section is to investigate which dimensions of heterogeneity that affect the MPC out of the housing wealth shock. Based on theoretical predictions and previous empirical studies we focus on credit- and liquidity constraints, as well as age, as important dimensions of heterogeneity (see, e.g., Aladangady, 2017; Jappelli and Pistaferri, 2010; Campbell and Cocco, 2007; Disney et al., 2010; Browning et al., 2013).

The pure wealth effect (i.e., how the lifetime wealth of homeowner changes with home prices) has been argued to likely be small (Buiter, 2010). In theory, the housing wealth can only contribute to consumption through a sale that produces a positive net value. This value is then discounted to the present, leading to consumption smoothing in the case of no credit constraints. For older homeowners, consumption is then smoothed out over fewer periods, leading to a greater MPC if it is mainly due to the pure wealth effect.

The sign of the wealth effect depends on the *expected* future net sale (Fagereng et al., 2024). However, a home sale is often in conjunction with a home purchase, which reduces the net value of the transaction, and negative in the case of upsizing. Only at downsizing events or moves to cheaper housing markets, can value be extracted. On the other hand, for renters and young households who expect to buy more housing in the future, changes in home prices affect their perceived lifetime wealth.

The collateral effect works through moving the borrowing constraint of homeowners who use their homes as collateral for mortgages. In the case of credit constrained homeowners, small moves in home prices can generate great shifts in the nominal cap on credit. These households, as well as liquidity constrained households, exhibit greater MPCs out of transitory income shocks as well as home price shocks (Parker et al., 2013; Lustig and Van Nieuwerburgh, 2005; Berger et al., 2018). High indebtedness correlates with being younger, and thus, if the collateral channel is stronger than the pure wealth effect, we expect younger homeowners to respond stronger than older.

---

our main regression sample.

Our procedure for disentangling these channels is to estimate differences in MPCs by household types. We construct dummy variables that split our sample population into several groups and include these dummy variables as an interaction term with our main explanatory variable. That is, in the second stage, we estimate for each consumption category (we will suppress the indicator for the consumption category),

$$\begin{aligned} c^k_{i,2014} - c^k_{i,2012} \quad &= \sum_{g\in\mathcal{G}} \beta^k_{g,\Delta t}\, \widehat{\Delta hw}_{i,2012+\Delta t} \times z_{g(i),2013} \\ &\quad + \eta^k_2\, X_{i,2012} + \alpha^k_2 + \varepsilon^k_{2\,i,\Delta t}. \end{aligned} \tag{6}$$

Compared with (5) the difference is thus the inclusion of $z_{g(i),2013}$s.[7] In our preferred specifications, $z_{g(i),2013}$, are time-invariant dummy variables splitting the households into equally-sized groups based on the value of the heterogeneity variable in 2013, that is, before the housing shock.

We construct four dummy variables and run regression models with each of them. Two measures proxy for credit constraints - debt-to-income (DTI) and loan-to-value (LTV). The other two are age (of oldest household member), and liquidity, measured as the total bank deposits. We then assign each household to one group.[8]

For LTV, we split up the sample in intervals of 10%, i.e., the group with the lowest LTV has a value in the range of 0–0.1, the second-lowest group is in the range of 0.1–0.2, and so forth. DTI divides households in quintiles, age by 6 years, and LTI by units of 2.5.

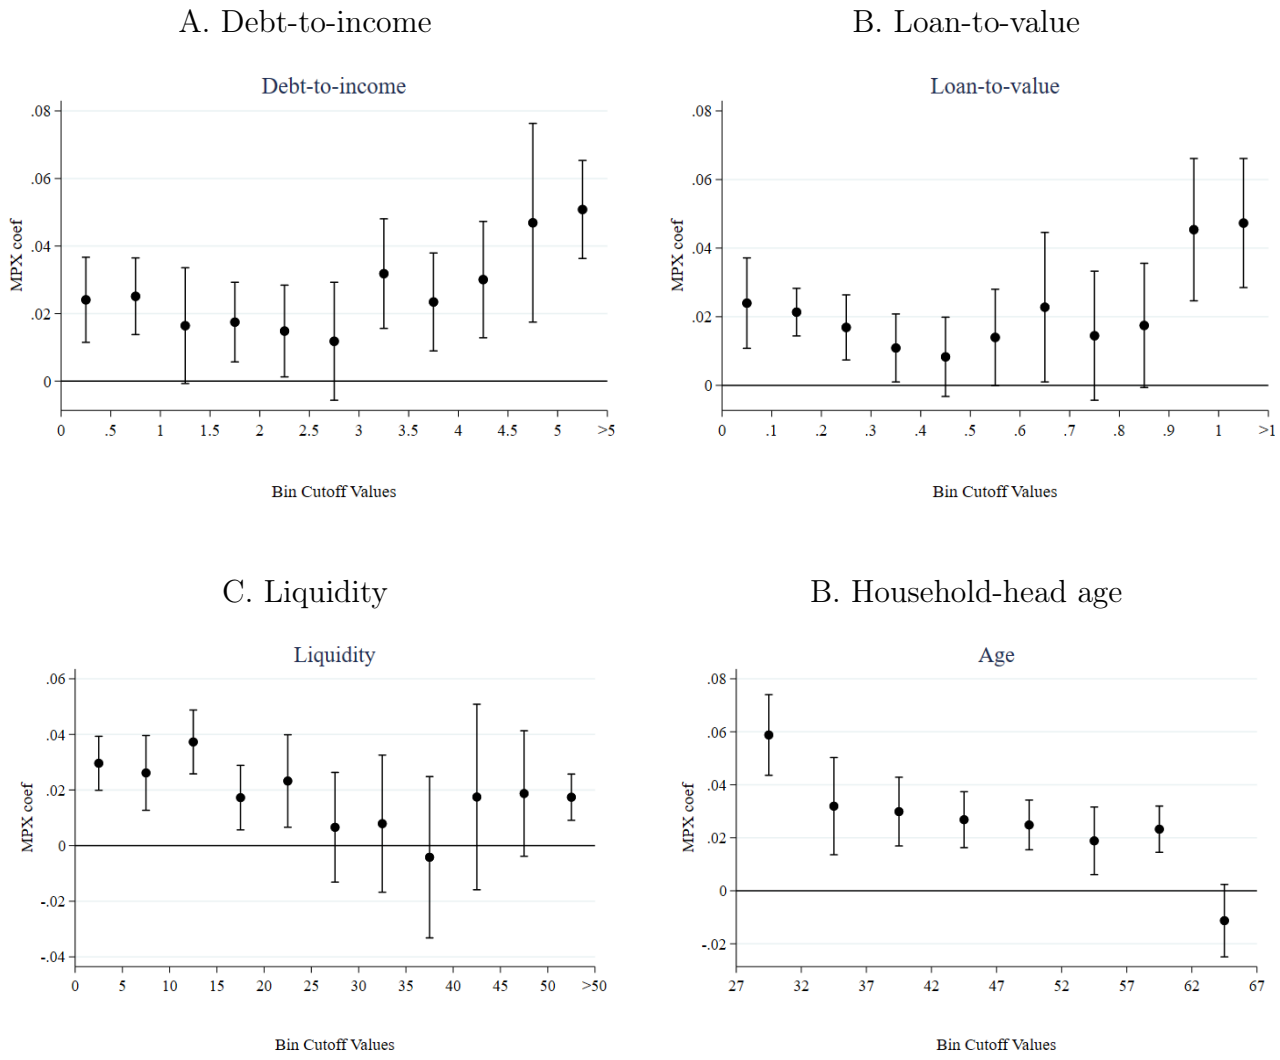
The estimates of (6) are presented in Figure 8 and display rich heterogeneity in the MPX. First, as debt-to-income rises, so does the MPX. For homeowners with the smallest DTI, the effect is zero, but rises almost linearly with DTI. For the most indebted group, the total MPX is 0.04, twice that of the baseline. Also for LTV we see an increasing effect. The effect

[7]Note that, since $z_{g(i),2013}$ is time-invariant, it is absorbed by the household fixed effects and should therefore be excluded from the regression.

[8]Different measurements and splits of the samples produce overall the same qualitative conclusions. Recall that, in Norway, since 2010 borrower-based requirements have limited households' maximum DTI and LTV. In 2013 these requirements were DTI less than five times gross income, and LTV to less than 80 percent.

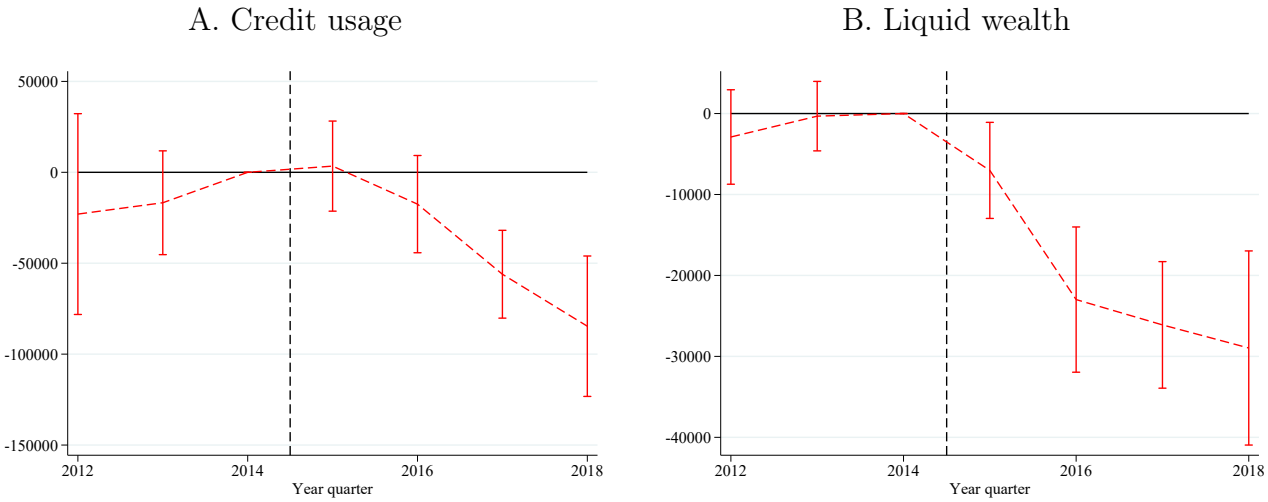
peaks at an LTV of 85%, coinciding with the legal cap on mortgages. After this point, the effect is if anything diminishing.

Figure 8: Heterogeneity in the MPXs



Note: This figure shows the heterogeneity of the total MPX for different household characteristics estimated using (6), while including household controls. Errors are cluster-robust at the level of the housing market and the confidence intervals are of 95%.

Figure 9: The evolution of credit and liquid wealth



Note: This figure shows amount of credit affected and non-affected households use, as well as the evolution of their liquid assets, estimated using eq. (1).

Liquidity displays the opposite relationship, in line with the importance of credit: as liquidity increases, the effect falls. For very high values the effect seems to rise but it is imprecisely estimated. The heterogeneity in age is in line with our expectations as the MPX falls consistently as consumers age. All these results are in line with an important role for housing as collateral to finance consumption. Homeowners who are likely credit constrained are responding the most while older homeowners do not respond—i.e., the pure wealth effect is small.

We provide further evidence that credit plays an important role in understanding the heterogeneity. We estimate a model like (1) but replacing the post_t dummy with year dummies (excluding 2014). The estimates are presented in Figure **??** and exhibit a significant drop in the increase in debt following 2014, albeit, no significant effect in 2015.

4.3 Summing up

Benmelech et al. (2022) analyze homebuyers and estimate that their expenditures on home durables and improvements triple in a two year period around their home purchase. As home purchases declined steeply during the Great Recession, the resulting drop in these expenditures likely played a substantial role in the propagation of the downturn. We disclose a similar mechanism for a much broader group of households, i.e., home owners react in a similar, negative, way when house prices fall by cutting back on expenditures on home durables and improvements. TBD

5 Conclusions

Understanding how shocks to housing wealth influence homeowners' expenditures is important for both policymakers given its implications for economic stability and business cycle fluctuations. Empirical estimates of the marginal propensity to consume out of housing wealth range from about three and seven percent. Particularly studies on the period around the Great Financial Crisis find numbers in the upper range, while studies using other instruments for housing price shocks indicate somewhat lower MPCs. Differences in accompanying economic conditions, such as stock market declines or a depressed labor market have been suggested as as explanations Little is yet known about the dynamics of consumption responses, heterogeneity across households, or the decomposition of dynamic consumption responses.

We present new evidence using the near universe of electronic payments for the full Norwegian population. We exploit a quasi-experimental setting around the 2014–2015 oil price collapse that caused local declines in home prices in the Stavanger region where the Norwegian oil industry is concentrated. The granularity and high frequency nature of the data enable us to map the full dynamic consumption response of households to the change in housing wealth and to analyze the plausibility of the aggregate response by breaking it

down into standardized expenditure categories.

In our analysis we compare government workers in the Stavanger region - the region in Norway where oil production is concentrated - who are exposed to the house price decline but face no change in income prospects due to centrally organized wage setting, with similar workers in the rest of Norway, Our analysis indicates that the average MPC out of housing wealth is approximately 0.02, somewhat smaller than what has been found in other studies, with significant variation across households and consumption categories.

Our findings demonstrate that the fall in home prices particularly led to reduced spending on durable goods such as furniture and vehicles, while essential consumption categories such as food and beverages remained largely unaffected. Looking into heterogeneity of the consumption response, we find strong evidence that the housing wealth effect on consumption is more pronounced in areas where credit constraints are tighter, particularly for homeowners with higher debt-to-income ratios, higher loan-to-value ratios, and lower liquidity-to-income ratios. The heterogeneity in consumption responses underscores the significance of the collateral channel over the wealth channel, particularly for credit-financed, durable goods.

The high frequency of our payments data reveal that the commonly observed cumulative MPX masks a substantial non-linearity of the MPX over time. At the start of the house price lapse consumption responds instantly while price effects initially are small, creating high MPXs. As price effects accrue, the contemporary MPX declines before it starts rising again with the cumulative house price effect building up.

These results suggest that while fluctuations in home prices do affect consumption, the magnitude of this effect may be smaller than previously estimated, especially in the absence of broader macroeconomic shocks like stock market declines. Moreover, the impact of reduced spending on durable goods will have limited implications for the local economy when durable goods are not locally produced.

Our findings also have implications for the design of policies aimed at mitigating the effects of house price shocks as MPCs are mainly high among households with high leverage

and low liquidity.

References

Aladangady, A. (2017). Housing wealth and consumption: evidence from geographically linked microdata. American Economic Review 107(11), 3415–3446. 1, 4, 4.2

Andersen, A. L., N. Johannesen, and A. Sheridan (2024). Dynamic Spending Responses to Wealth Shocks: Evidence from Quasi-lotteries on the Stock Market. American Economic Review Insights (Forthcoming). 1

Andersen, H. Y. and S. Leth-Petersen (2020, 03). Housing Wealth or Collateral: How Home Value Shocks Drive Home Equity Extraction and Spending. Journal of the European Economic Association 19(1), 403–440. 1

Aruoba, S. B., R. Elul, and Kalemli-Özcan (2022, October). Housing wealth and consumption: The role of heterogeneous credit constraints. Working Paper 30591, National Bureau of Economic Research. 1

Attanasio, O. P., L. Blow, R. Hamilton, and A. Leicester (2009). Booms and busts: Consumption, house prices and expectations. Economica 76(301), 20–50. 4.1

Baker, R. S., R. A. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis (2023, 03). Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments*. Review of Finance 27(6), 2271–2304. 1

Benmelech, E., A. Guren, and B. T. Melzer (2022, 07). Making the House a Home: The Stimulative Effect of Home Purchases on Consumption and Investment. The Review of Financial Studies 36(1), 122–154. 1, 4.3

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2018). House prices and consumer spending. The Review of Economic Studies 85(3), 1502–1542. 1, 4.2

Berger, D. and J. Vavra (2015). Consumption dynamics during recessions. Econometrica 83(1), 101–154. 1

Böjeryd, J., D. Kessel, B. Tyrefors, and R. Vestman (2023). The housing wealth effect: Quasi-experimental evidence. CEPR Discussion Paper DP18034. 1, 4.1.1

Browning, M., M. Gørtz, and S. Leth-Petersen (2013). Housing wealth and consumption: a micro panel study. The Economic Journal 123(568), 401–428. 1, 4.2

Buiter, W. H. (2010). Housing wealth isn't wealth. Economics: The Open-Access, Open-Assessment E-Journal 4. 4.2

Campbell, J. Y. and J. F. Cocco (2007). How do house prices affect consumption? evidence from micro data. Journal of Monetary Economics 54, 591–621. 1, 4.2

Carroll, C. D., M. Otsuka, and J. Slacalek (2011). How large are housing and financial wealth effects? a new approach. Journal of Money, Credit and Banking 43(1), 55–79. 1

Disney, R., J. Gathergood, and A. Henley (2010, December). House price shocks, negative equity, and household consumption in the united kingdom. Journal of European Economic Association 8(6), 1179–1207. 4.1.1, 4.2

Eggertsson, G. B. and P. Krugman (2012). Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach. The Quarterly Journal of Economics 127(3), 1469–1513. 1

Fagereng, A., M. Gomez, E. Gouin-Bonenfant, M. Holm, B. Moll, and G. Natvik (2024). Asset-price redistribution. Working paper. 4.2

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

Graham, J. and C. A. Makridis (2023, January). House prices and consumption: A new instrumental variables approach. American Economic Journal: Macroeconomics 15(1), 411–43. 1

Guerrieri, V., G. Lorenzoni, and M. Prato (2020, 11). Schumpeter Lecture 2019: Slow Household Deleveraging. Journal of the European Economic Association 18(6), 2755–2775. 1

Guren, A., A. McKay, E. Nakamura, and J. Steinsson (2021). Housing wealth effects: The long view. The Review of Economic Studies 88(2), 669–707. 1, 4.1.1

Imas, A., T.-C. Lin, Y. Luo, and X. Wang (2022). In victory or defeat: Consumption responses to wealth shocks. Working Paper, University of Chicago, Becker Friedman Institute for Economics. 1

Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. Annual Review of Economics 2(1), 479–506. 4.2

Kaplan, G., K. Mitman, and G. L. Violante (2020). Non-durable consumption and housing net worth in the great recession: Evidence from easily accessible data. Journal of Public Economics 189, 104176. 1

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. Econometrica 82(4), 1199–1239. 1

Laibson, D., P. Maxted, and B. Moll (2022, January). A Simple Mapping from MPCs to MPXs. NBER Working Papers 29664, National Bureau of Economic Research, Inc. 3

Lustig, H. N. and S. G. Van Nieuwerburgh (2005). Housing collateral, consumption insurance, and risk premia: An empirical perspective. The Journal of Finance 60(3), 1167–1219. 4.2

Meyer, S., M. Pagel, and A. Previtero (2020). The consumption response to realized capital gains: Evidence from mutual fund liquidations. Mimeo. 1

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. 1

Mian, A. and A. Sufi (2014). What explains the 2007–2009 drop in employment? Econometrica 82(6), 2197–2223. 1

Paiella, M. and L. Pistaferri (2017, 10). Decomposing the Wealth Effect on Consumption. The Review of Economics and Statistics 99(4), 710–721. 1

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. 4.2

Sinai, T. and N. S. Souleles (2005). Owner-Occupied Housing as a Hedge Against Rent Risk*. The Quarterly Journal of Economics 120(2), 763–789. 1

United Nations (2018). Classification of Individual Consumption According to Purpose (COICOP). Statistical Papers Series M, No.99, UN Statistics Division. 2, 2.1

Appendices

1 Government workers

We define a worker as employed by government if their "institutional sector" (*institusjonell sektorgruppering*) is 1110 or 1120 (non-financial enterprises owned by central government), 1510 or 1520 (non-financial enterprises owned by local government), 3100 or 3900 (public financial corporations), 6100 (central government), or 6500 (local government) when using the version from 2012. The codes we use for older observations use the 1987 version and are 110 (central government) 510 (county municipalities), 550 (municipalities), 150, or 190 (public financial corporations), 610, 630, 635, 660, or 680 (public, non-financial enterprises). We exclude workers in industries (Standard Industrial Classification 2007) that start with 06 (extraction of crude petroleum and natural gas) and 09.1 (support activities for petroleum and natural gas extraction).