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

Knut Are Aastveit[†] Jesper Böjeryd[‡] Magnus A. H. Gulbrandsen[§]

Ragnar Enger Juelsrud[¶] Kasper Roszbach[‖]

November 2025

Abstract

We study how changes in housing wealth affect household spending using administrative and granular, de-identified, data on debit card payments and e-invoices for the near population of Norway. We focus on the 2014 oil-price collapse, which created sharp regional variation in house prices. Comparing government workers in oil and non-oil regions, we estimate a three-year marginal propensity to spend (MPX) of about 3.6 cents per dollar. The response is highly concentrated in durables, home improvements, furnishings, and vehicles, and primarily driven by a reduction in the uptake of credit backed by home-equity. The local MPX (L-MPX), the share of the total spending response accruing to locally produced goods and services, stands for roughly 80% of the consumption decline. Additional findings highlight both collateral and wealth effects as key channels linking housing wealth to consumption, and document that household balance-sheet heterogeneity shapes the propagation of housing-wealth shocks.

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

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

*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 thank Yann Cerasi and Elias Brødholt for excellent research assistance. We are grateful to our discussants Adi Aladangady, Denis Gorea, Matthijs Korevaar, Niels Johannesen, and Vincenzo Pezone and to Dominic Cucic, Joao Cocco, Daria Finocchiaro, Martin Flodén, Andreas Fuster, Refet Gurkaynak, Steinar Holden, Thomas Jansson, Benjamin Keys, Karin Kinnerud, Søren Leth-Petersen, Jonas Niemeyer, Pascal Paul, Federico Ravenna, Yannick Timmer, Neeltje van Horen, Karl Walentin, and other seminar and conference participants for valuable comments. We also thank participants at the ABN AMRO Workshop on Pioneering Economics with Financial Transaction Data, Bank of Greece, Berlin Macro Seminar at DIW, BI Norwegian Business School, CEBRA Annual Meeting in Frankfurt, Central Bank of Ireland, Czech National Bank Workshop on Financial Stability and Macroprudential Policy, Federal Reserve Bank of Atlanta, IAAE Conference in Turin, the NV FEB workshop at Erasmus University, SNDE Conference in San Antonio, Statistics Norway, Stockholm University, Sveriges Riksbank, Swiss National Bank, Swiss Winter Conference on Financial Intermediation, Tilburg Finance Summit, Universitat Autònoma de Barcelona, and the conference on *Monetary Policy and Financial Intermediation: Learning from Heterogeneity and Microdata* at Collegio Carlo Alberto, and the University of Southern California. This paper is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at BI Norwegian Business School.

[†]Norges Bank & BI Norwegian Business School. Email: knut-are.aastveit@norges-bank.no

[‡]Stockholm Universtity. Email: jesper.bojeryd@su.se

[§]Norges Bank, Email: magnus.gulbrandsen@norges-bank.no

[¶]Norges Bank, BI Norwegian Business School & CEPR. Email: ragnar.juelsrud@norges-bank.no

[‖]Norges Bank & University of Groningen. Email: kasper.roszbach@norges-bank.no

1 Introduction

How do changes in housing wealth affect household consumption? Housing is the largest asset on household balance sheets and the primary form of collateral for borrowing. Fluctuations in house prices therefore have the potential to shape household demand, propagate shocks across regions, and influence the transmission of monetary and macroprudential policy. A large literature documents that house prices and consumption move together, emphasizing both collateral constraints and wealth effects as drivers (Mian et al. (2013); Mian and Sufi (2014); Aladangady (2017); Berger et al. (2018).[1] Despite the importance of the topic, key questions remain. What is the relative role of the wealth and collateral channel in explaining the link between housing wealth and consumption? Are consumption responses broad-based or tied to particular spending categories, such as housing related expenditures, and to what extent do they affect demand for locally produced goods and services? Both questions are central for understanding how to model the link between house prices and consumption, and for assessing the importance of the housing market for macroeconomic policy. Yet, due to the limited availability of granular consumption data, little is still known on these two fronts. This paper aims to bridge this gap.

We answer these questions by combining uniquely detailed, de-identified, Norwegian data on actual spending with a clean source of exogenous variation in housing wealth. Our data link administrative balance-sheet information and income records to transaction-level card and e-invoice spending for nearly the entire population over the period 2012-2018, allowing us to observe actual expenditures at high frequency and across 20 detailed consumption categories. This granularity enables us to quantify how housing-wealth shocks affect both the level and composition of household spending - something that earlier studies could only

[1] U.S. studies often find large effects of 5–7 cents, consistent with powerful collateral channels (Mian et al., 2013; Aladangady, 2017). More recent work such as Guren et al. (2021) argues for smaller effects of about three cents once general equilibrium spillovers are taken into account. European registry-based studies typically find more modest responses, around 2–3 cents (Browning et al., 2013; Vestman et al., 2023; Paiella and Pistaferri, 2017). At the other end of the spectrum, models with incomplete markets and collateralized borrowing predict larger responses, with calibrations in Berger et al. (2018) and Kaplan et al. (2020) implying effects approaching ten cents under certain conditions.

approximate. While earlier registry-based research estimates total spending responses, we quantify how much of this response accrues to locally produced goods and services, allowing for a new mapping between micro-level wealth effects and local demand multipliers. Our empirical setting is Norway, where the homeownership rate was and is approximately 80 percent, during the 2014 oil-price collapse that triggered large regional variation in house prices due to the geographic concentration of oil-sector employment. We compare home-owning *government workers* in "oil regions" and "non-oil regions", who experienced similar income dynamics but divergent housing-wealth trajectories. Because public-sector wages are centrally negotiated and nationally uniform, this comparison isolates housing-wealth changes from income or employment risk.

We find that the average marginal propensity to spend (MPX) out of housing wealth is around 3.6 cents per dollar over a three-year horizon. The spending response is highly concentrated in a few durable consumption categories - home improvements, furnishings, and vehicles in particular. Because we have access to highly disaggregated expenditures, we proceed by introducing the concept of a "local marginal propensity to spend (L-MPX)" that quantifies the share of the total spending response accruing to locally produced goods and services. When we accounting for import leakages, the local demand impulse is about 80 percent of the headline estimate, suggesting that the effects of housing wealth changes on local demand can be substantially smaller than the headline estimates.

Beyond the aggregate response, we examine the mechanisms linking housing wealth to spending. Our evidence points to home-equity extraction (HEE) as the dominant transmission channel: households facing tighter collateral constraints respond nearly twice as strongly as unconstrained ones. This pattern indicates that housing wealth affects spending not only through broad revaluation effects, but also through borrowing constraints that amplify downturns. At the same time, we document meaningful wealth effects even among unconstrained households, implying that both channels operate in parallel. Taken together, these results show that the aggregate spending response to housing wealth depends crucially

on the distribution of debt, liquidity, and access to refinancing across households.

We organize the paper into five complementary steps. First, we link administrative records on the near universe of Norwegian households, including detailed balance sheets and income statements, to 12 billion debit card and e-invoice transactions across 26 COICOP consumption categories, an international classification system to categorize household consumption expenditures, at the household-zip-week level.[2] Our spending data capture most of household consumption in Norway; compared to official survey data from Statistics Norway, our aggregated microdata comprise about 91% of aggregate consumption levels and closely tracks the dynamics of consumption in the national accounts.

Second, we exploit a quasi-natural experiment in Norway that isolates changes in housing wealth from changes in income or employment risk during the 2014 oil-price collapse, and focus on homeowning government workers living inside and outside these "oil regions". While this price shock triggered sharp regional house-price declines concentrated in petroleum-exposed areas, public-sector wages and employment were left unaffected because they are centrally negotiated and nationally uniform. Our research design therefore provides a clean comparison of households with similar and stable incomes but diverging housing-wealth dynamics. We implement this comparison in an instrumental-variables (IV) framework, using residency in the oil region as an instrument for changes in housing wealth over the 2014–2017 period. With government workers in oil regions experiencing roughly a 25 percent larger decline in housing wealth than those elsewhere, our instrument is relevant and strong. It is also plausibly exogenous to other determinants of consumption. This design allows us to identify the causal effect of housing wealth on spending with high internal validity.

Third, we estimate a marginal propensity to spend (MPX) out of housing wealth of roughly 3.6 cents per dollar over a three-year horizon, identified using regional variation from the 2014 oil-price collapse. The estimate is robust to a wide range of validity checks. Government workers in oil and non-oil regions display parallel spending trends before the

[2]COICOP: United Nations Classification of Individual Consumption by Purpose (United Nations, 2018).

shock, and after the collapse, the groups experience similar income and capital-gains trajectories. The response is also unrelated to local unemployment dynamics, i.e., peer or regional spillovers do not drive our findings. Taken together, these results provide a credible and precise estimate of the average MPX out of housing wealth, situated near the middle of estimates in the international literature.

Fourth, we exploit the granularity of our transaction data to decompose the spending response across 20 detailed consumption categories. The effects are highly concentrated in a few durable consumption categories - home improvements, furnishings, and vehicles - typical purchases to be financed through home equity withdrawals. This concentration implies that the overall spending response to housing-wealth changes is not broad-based but driven by goods and services that are sensitive to the availability of collateral-backed credit. From a policy perspective, this distinction is important because macroprudential concerns about the effects of house-price declines typically center on local demand spillovers and on how a collateral channel may amplify regional downturns. To quantify this, we introduce the concept of a local marginal propensity to spend (L-MPX), which measures the portion of the aggregate response accruing to locally produced goods and services. We derive sector-specific MPXs from our category estimates and combine them with input–output tables to weight each sector by its local production share. The resulting L-MPX is about 80 percent of the headline MPX, indicating substantial leakage through imports and inter-regional trade. This distinction between total and local spending responses connects our findings to Farhi and Werning (2016), who show that trade and financial linkages attenuate local multipliers, and to Guerrieri et al. (2020), who highlight that slow household deleveraging can prolong recessions even without income shocks. Together, these results demonstrate that both the magnitude and the composition of spending responses are central to understanding how housing-wealth shocks propagate through the economy.

Fifth, we examine the mechanisms through which housing-wealth changes affect household spending. Our evidence reveals strong heterogeneity by financial position. The MPX

rises sharply with household leverage - those with high loan-to-value or debt-to-income ratios - reaching six to seven cents per dollar. Among unconstrained households, i.e., with low leverage or ample liquidity, the wealth effect is still sizeable, about 2.5 cents per dollar, and consistent with a pure revaluation channel. Home-equity extraction (HEE) emerges as the central conduit linking housing-wealth declines to spending. Government workers in affected oil regions curtailed new borrowing, drew down deposits, and sharply reduced expenditures on home improvements—the single category that fully explains the MPX gap between constrained and unconstrained households. This finding complements the insights from Benmelech et al. (2022), who show that buying a house (the extensive margin) stimulates expenditures on home improvements. We instead document that home improvement expenditures are an important also on the intensive margin, through house price fluctuations. In our setting, falling prices tighten borrowing constraints, suppress home improvements, and amplify local downturns. Together, these findings demonstrate that collateralized borrowing and HEE are not peripheral details but make up the core transmission mechanisms of housing-wealth shocks, and that macroeconomic models will need to incorporate these to capture the interplay between household balance sheets and aggregate demand.

Relation to the Literature Our paper contributes to several strands of the literature. The first one empirically examines the link between housing wealth and consumption, emphasizing different magnitudes and mechanisms. U.S. evidence by Mian et al. (2013), Mian and Sufi (2014) and Aladangady (2017) highlights a collateral channel: regions with larger house-price declines experienced sharper spending contractions, particularly among leveraged households. Other work emphasizes pure wealth effects that are independent of collateral constraints. Berger et al. (2018) show that even unconstrained households adjust spending when house values change, with the magnitude depending on the distribution of debt and the persistence of shocks. Complementary studies including Sinai and Souleles (2005), Campbell and Cocco (2007), and Carroll et al. (2011) emphasize life-cycle and buffer-stock

mechanisms that generate wealth effects even without borrowing constraints. Together, this U.S. evidence suggests that both collateral and wealth channels matter, with relative importance shaped by leverage, liquidity, and expectations. Graham and Makridis (2023) introduce a new instruments for U.S. house prices and find smaller marginal propensities to spend for non-durables, about 0.8–0.9 cents per dollar. Evidence from Europe also points to smaller aggregate effects of housing wealth on spending (Browning et al., 2013; Vestman et al., 2023; Paiella and Pistaferri, 2017), likely reflecting differences in institutional features of mortgage markets, such as stricter refinancing rules and limited access to HEE, rather than the absence of wealth effects. Andersen and Leth-Petersen (2020) identify unanticipated home-value changes in Denmark and show that both consumption and home equity extraction respond. Benmelech et al. (2022) highlight the importance of house purchases for consumer spending on home improvements. We add to this literature by documenting the granular components of the consumption response to a drop in house prices, and establish that home ownership and home equity are important vehicles for the transmission of shocks on the intensive margin as constrained households reduce borrowing and scale back on home improvements. The very channel that boosts local demand in booms can thus generate or reinforce contractions during busts.

A second strand of research studies the nature and persistence of consumption responses to other forms of wealth and income shocks than house prices. Andersen et al. (2024) exploit stock-market lottery windfalls in Denmark and find strong but short-lived spending effects, underscoring that not all wealth shocks operate similarly. Heterogeneous responses also emerge in studies of fiscal and labor-market shocks (Meyer et al., 2020; Berger and Vavra, 2015; Fagereng et al., 2021; Chodorow-Reich et al., 2021; Baker et al., 2023; Imas et al., 2022), reinforcing that the marginal propensity to consume depends critically on the liquidity of the shock and household balance-sheet positions. Some studies connect the housing-wealth effects to macroeconomic transmission and local multipliers. Beraja et al. (2019) and Guren et al. (2021) emphasize that regional consumption multipliers are attenuated when spillovers

operate through trade, migration, or financial linkages. Farhi and Werning (2016), Auclert et al. (2020), and Guerrieri et al. (2020) provide theoretical frameworks linking household balance sheets and incomplete markets to the aggregate propagation of regional shocks. We add to this literature by providing new, direct, microdata-based evidence on how much of the spending response to housing wealth accrues to locally produced goods. We refer to this new measure as the *local marginal propensity to spend* (L-MPX). This connects household-level wealth effects to the geography of demand transmission, and links micro evidence on collateral and borrowing constraints to macro models of regional multipliers and aggregate propagation.

Finally, our findings also relate to models of consumption under incomplete markets and collateralized borrowing. Kaplan and Violante (2014) and Kaplan et al. (2020) show that heterogeneity in debt, liquidity, and expectations generates variation in spending sensitivities that are consistent with our evidence. Aruoba et al. (2022) use auto loan origination data to proxy for individual-level consumption and show, for the US, that a large part of the consumption response that earlier research had attributed to wealth effects can instead be explained by heterogeneity in households' financial constraints. Our results extend this literature by exploiting highly granular register-based expenditure data and show that both collateral and wealth effects act as key channels that link housing wealth to consumption. We additionally document that household balance-sheet heterogeneity plays a an important role in quantitatively shaping the propagation of housing wealth shocks.

Taken together, this literature underscores that the magnitude and persistence of housing-wealth effects depend on factors such as balance-sheet heterogeneity and the institutional context. Our contribution is to document at a unique level how granular components of consumption respond to housing wealth shock. We quantify how collateral constraints and home-equity extraction shape household spending and to demonstrate how these mechanisms determine both the composition and the *localness* of demand. By introducing and quantifying a local marginal propensity to spend, the L-MPX, we provide a bridge from micro-level

evidence on household balance sheets to macro models of regional spillovers and aggregate propagation, and contribute to a better empirical foundations to understand how housing-wealth shocks transmit through the real economy.

The remainder of the paper is organized as follows. Section 2 describes the data; Section 3 outlines the empirical design; Section 4 presents aggregate and disaggregated results and introduces the L-MPX; Section 5 examines mechanisms; and Section 6 concludes.

2 Data and sample

This section describes our data sources, the construction of the regression sample, and the definition of key variables. We then present summary statistics that compare treatment and control groups to validate our design.

2.1 Data sources

We combine four data sources. The first is expenditure data for the full adult Norwegian population at the individual × week × zip code level. The data are provided by NETS, which operates the technical infrastructure for in-store and online payments in Scandinavia, the Baltics, and Switzerland. NETS services BankAxept, a company jointly owned by the largest Norwegian banks that hosts an open system for debit and electronic payments. Nearly all non-financial firms are connected, and any bank operating in Norway can join. BankAxept processes debit card transactions, electronic bill payments, direct remittances, and cash withdrawals for the universe of Norwegian residents and banks, and accounted for approximately 71%[3] of total card transaction volume over our sample period. All electronic invoice payments and direct remittances run through banks connected to BankAxept.

To protect privacy, expenditures are aggregated from individual payments into 26 categories based on COICOP, the United Nations classification for consumption expenditures

[3]Source: Norges Bank, Payments Statistics.

(United Nations, 2018), and we further aggregate somewhat to 20 consumption groups. Our data therefore provides high-frequent household spending data across 20 consumption groups, including large durables such as vehicles and furnishings that are typically underrepresented in survey data. Because credit-card spending is repaid through e-invoices, we observe total card outlays at the bill level but not their category breakdown. Categories such as foreign hotel bookings may therefore be underrepresented, though these are small in Norway given the prevalence of debit-card payments.

The second source is administrative tax and employment registers from Statistics Norway (SSB). Norway levies both income and wealth taxes, so registers provide a comprehensive view of individual balance sheets and income sources (deposits, debt, housing wealth; wage, business, and capital income). Identifiers allows us to aggregate from individuals to households. Third, we use transaction-level real-estate records for all Norwegian housing sales, including transaction prices and property characteristics. We complement this with a fourth dataset on regional house-price indices from Eiendom Norge to measure time-series variation in local house values. Together, these sources allow us to construct household-specific housing wealth and link it to detailed spending and balance-sheet outcomes.

2.2 Sample construction

Our baseline sample consists of households in which strictly more than half of adult members worked in the government sector in both 2013 and 2014. The motivation for this focus is to ensure that households in our sample had limited changes in income prospects or income risks over this period. We revisit and discuss this issue at length in Section 4.3.

The initial definition yields 281,000 government-worker households. Restricting to homeowners at the start of 2012, excluding households with more than 5% self-employment income, and removing retirees and other special groups results in a final sample of approximately 117,000 households.

2.3 Variable definitions

In this subsection, we define the key variables used in the analysis. Unless otherwise noted, all monetary amounts are nominal Norwegian Krone (NOK).[4]

Consumption Our main outcome is household consumption. In contrast to national accounts, we use "consumption", "expenditures" and "spending" interchangagebly. We define total consumption as the sum across all categories, net of payments to the public sector and other non-classifiable expenditures. To remove outliers, we exclude transactions related to housing purchases, vehicle transactions above NOK 5,000,000, and other transactions above NOK 500,000.

For heterogeneity analysis, we use 20 disaggregated categories to study the composition of spending responses and to distinguish margins likely to have local spillovers.

A relevant question is how much coverage our consumption measure in terms of total consumption. To benchmark our consumption measure against the national accounts, we therefore compare aggregate expenditures in our data with Statistics Norway (SSB) consumer surveys. Table 1 shows that our expenditures are slightly below survey figures overall (within 1.8–15% depending on category). Differences largely reflect category definitions (e.g., the survey includes individual-to-individual used vehicle purchases, which are excluded in our debit/e-invoice measure).

An important advantage of our data is its high frequency. To smooth within-household seasonality and reduce high-frequency noise, we compute a moving average of each household's consumption over the current and three subsequent quarters:

$$\bar{c}_{i,t} \equiv \frac{c_{i,t} + c_{i,t+1} + c_{i,t+2} + c_{i,t+3}}{4}. \tag{1}$$

This yields a smoother series that preserves medium-run fluctuations relevant for estimating marginal propensities to spend (MPX).

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

Table 1: Comparison to consumer surveys

	Total consumption	Vehicles	Furnishing	Food & beverages
Microdata	397,776	15,428	24,936	91,880
Survey data:				
All Norway	435,507	43,226	24,495	63,146
Rural	404,862	57,742	24,224	65,185
Urban	445,068	42,282	25,962	64,175
Agder & Rogaland	440,794	44,203	26,950	64,417
Two-person households	466,821	50,358	27,694	69,030

Note: This table compares our expenditure measure to survey data used in the construction of the national accounts. Survey categories do not coincide exactly with ours. "Vehicles" refers to COICOP group 071 (vehicle purchases) and includes used vehicles; "furnishing" refers to group 05 (furniture, household items, and maintenance); and "food and beverages" aggregates groups 01 and 02 (groceries, non-alcoholic beverages, alcohol, and tobacco). In the microdata, Food & beverages includes all spending in grocery stores. Urban is defined as non-rural municipalities excluding Oslo, Bergen, and Trondheim. We use the 2012 survey wave.

Household income and balance sheet variables Housing wealth is the imputed market value of primary and secondary residences, obtained by updating the beginning of sample (2012) tax value with local house-price indices. We measure leverage and liquidity using standard ratios:

- *Loan-to-value (LTV)*: total household debt divided by the imputed market value of primary and secondary residences.
- *Debt-to-income (DTI)*: total household debt divided by after-tax income.
- *Liquid-to-income (LTI)*: bank deposits divided by after-tax income.

These measures capture heterogeneity in exposure to collateral and liquidity constraints that is central to our interpretation of the results.

Demographics For each household, we compute the average age of tax-filing members, the number of adults employed in the government sector, the number of children under 18, and municipality of residence. These variables are used to assess sample comparability and as controls in regression analysis.

2.4 Summary statistics and sample balance

Government workers consists approximately 30% of the labor force. As a result, households in the government sector are fairly comparable to the rest of the population. Moreover, there is limited variation in the characteristics of government workers across regions. Tables 2 and 3 illustrate these points. Households in oil-exposed regions have somewhat higher housing wealth and debt, but differences are modest. Government-worker households are slightly older and less leveraged than private-sector households. Although we do not assume that households inside vs. outside oil region are balanced on observables, the fact that these two groups largely appear balanced ex ante reduces the risk that observed consumption differences reflect unobserved confounding factors rather than housing wealth.

Institutional context and external validity While our empirical setting is Norway, several institutional features make it broadly representative of advanced mortgage markets. Norwegian households exhibit high homeownership (around 80 percent), high leverage, and predominantly variable-rate mortgages. Refinancing and home-equity extraction are straightforward, similar to practices in the United States, but more flexible than in Denmark or Sweden. Mortgage interest deductibility and liquid secondary markets further support the comparability. Hence, although our identification leverages a Norwegian-specific shock, the mechanisms we document—wealth effects, collateral constraints, and equity extraction—are relevant for other advanced economies with developed mortgage markets.

Table 2: Summary Statistics – Governmernt workers in Oil vs. Non-oil regions

Variable	Oil region households			Non-oil region households		
	Mean	Median	SD	Mean	Median	SD
Panel A: Demographics						
HH head age	46.17	46.00	9.81	46.96	47.00	9.67
HH size	1.85	2.00	0.92	1.88	2.00	0.89
Panel B: Household income						
Labor inc.	781 413	680 208	363 690	790 953	705 458	363 794
Capital inc.	12 401	4 053	44 170	11 306	3 701	38 146
Unemp. inc.	1 147	0	13 817	2 099	0	17 586
Post tax inc.	617 832	580 508	258 091	627 203	604 394	254 910
Debt	1 835 123	1 699 500	1 315 722	1 635 354	1 489 000	1 158 378
Panel C: Other balance sheet						
Liquid wealth	307 392	151 858	469 983	303 991	142 642	471 482
Housing wealth	3 510 940	3 294 260	1 354 940	2 789 344	2 556 906	1 283 851
LTV	0.54	0.50	0.38	0.61	0.57	0.46
DTI	2.34	2.25	1.64	2.04	1.94	1.67
Panel D: Consumption						
Total cons.	415 166	354 393	275 234	410 365	359 082	260 784
Tradables	145 612	114 285	14 335	147 106	119 384	136 236
Non-tradables	159 058	130 096	135 391	151 388	127 030	122 753
Essentials	104 397	93 827	65 404	104 839	94 953	62 299
Non-essentials	180 365	133 383	185 922	173 582	133 035	171 208
Num. HHs	5 790	5 790	5 790	111 628	111 628	111 628

Note: This table contains summary statistics on key observable variables (listed in column one) for household classified as government worker households in the oil region (treatment group) and government worker households in the non-oil regions (control group). 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.

Table 3: Summary Statistics – Government-worker households vs. Private-sector households

Variable	Government-worker households			Private-sector households		
	Mean	Median	SD	Mean	Median	SD
Panel A: Demographics						
HH head age	46.92	47.00	9.68	44.11	44.00	9.59
HH size	1.88	2.00	0.90	2.17	2.00	0.89
Panel B: Household income						
Labor inc.	790,482	704,427	363,794	948,609	891,823	438,343
Capital inc.	11,360	3,716	38,466	19,827	4,502	73,928
Unemp. inc.	2,052	0	17,421	4,601	0	23,813
Post-tax inc.	626,741	603,383	255,076	750,823	723,855	298,164
Debt	1,645,204	1,496,000	1,167,434	2,104,931	1,902,000	1,437,244
Panel C: Other balance sheet						
Liquid wealth	304,158	143,091	471,409	319,039	159,249	503,941
Housing wealth	2,824,927	2,591,178	1,296,892	3,131,831	2,812,312	1,506,584
LTV	0.60	0.57	0.46	0.69	0.66	0.45
DTI	2.05	1.95	1.67	2.23	2.07	2.48
Panel D: Consumption						
Total cons.	410,602	358,852	261,517	513,587	460,039	311,524
Tradables	147,032	119,147	136,595	179,134	147,842	166,464
Non-tradables	151,766	127,181	123,418	187,316	163,446	148,956
Essentials	104,817	94,896	62,456	122,686	117,612	67,432
Non-essential	173,917	133,053	171,969	219,519	170,637	213,478
Num. hhs	117,418	117,418	117,418	420,741	420,741	420,741

Note: This table reports summary statistics on key observable variables (listed in column one) for households classified as government-worker households and private-sector households. A household is classified as a government-worker household if more than half of adult members were employed in the public sector in both 2013 and 2014. Private-sector households are defined analogously. All variables are measured in 2012, prior to the oil-price collapse. Income and wealth are annual and expressed in nominal NOK. Debt-to-income (DTI) is total household debt divided by after-tax income. Loan-to-value (LTV) is total debt divided by the imputed market value of housing. Liquid-to-income (LTI) is deposits divided by after-tax income. Government-worker households are slightly older, less leveraged, and consume less on average than private-sector households.

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.

3.1 Experimental setting

Our empirical strategy to identify the causal effect of house prices on consumption is to 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. In the ideal experimental setting, the econometrician would observe one-off and unpredictable variations in housing wealth at the household level. The quasi-experimental setting we exploit approximates the ideal conditions by providing us with household level variation in housing wealth induced by the sharp decline in the oil price in 2014.

Figure 1 illustrates the context of our empirical setting. In the left panel, we show the share of government workers in the local labor market relative to the average national share. As is clear from the figure, oil activities are highly concentrated in the south-west of Norway. In fact, most oil workers live in around the city of Stavanger, in the county of Rogaland. We refer to this area as the "oil region".

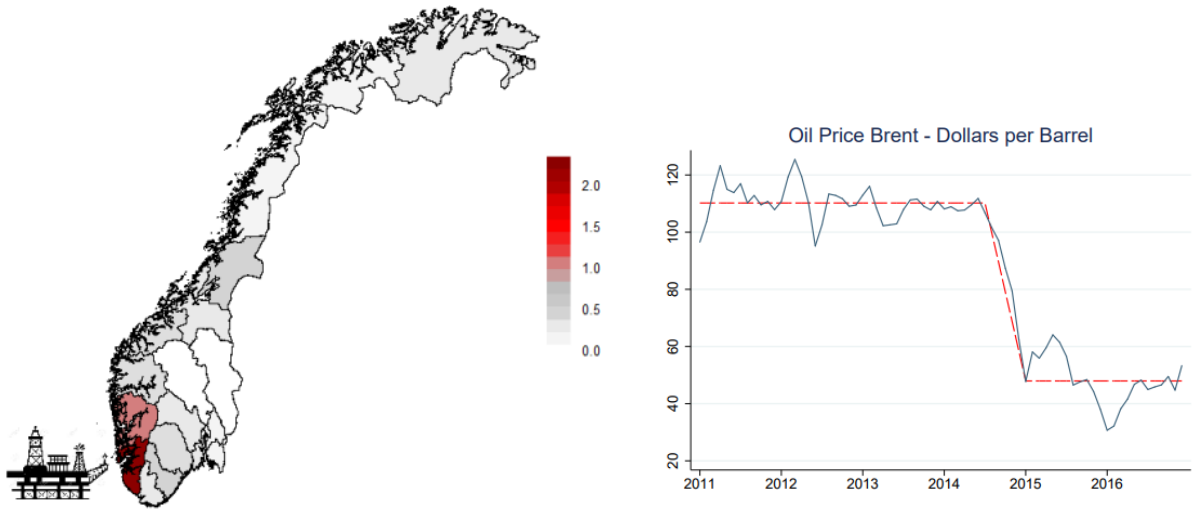
In the right panel, we show the evolution of the USD Brent Price. From 2011 to the beginning of 2014, the oil price was stable around USD 110. Then, over the course of half a year, the oil price fell down to about USD 50 and remained permanently lower. This large decline came as a surprise on market participants, as illustrated in future markets where oil futures for different medium-term horizons remained at 110 well into 2014.

When oil prices fell and stayed low, oil workers eventually relocated. While house prices were initially higher in the oil region compared to the rest of the country, over the course of 2015-2016 they started to fall and kept falling until the end of 2017, as illustrated in Figure 2.

Figure 1: Oil region and oil price

Panel A: Oil region

Panel B: Oil price (Brent USD)

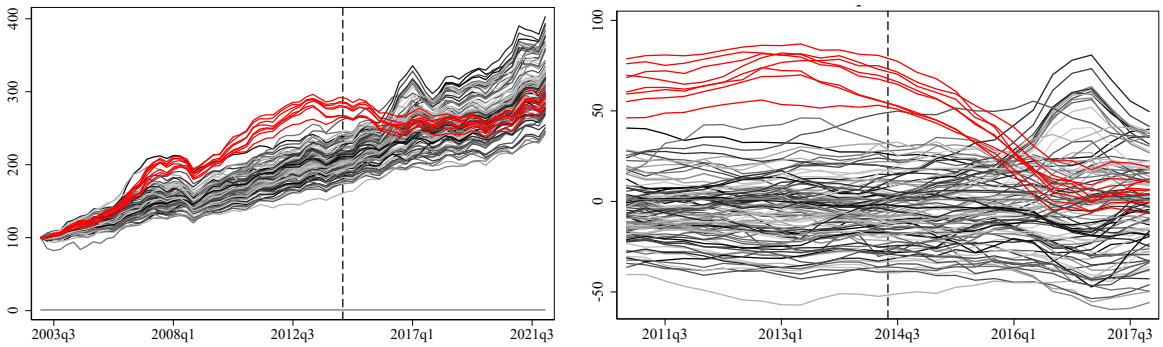


Note: The left panel displays the fraction of oil workers relative to the national average for different counties in Norway. The right panel displayes the Brent oil price per barrel. Horizontal lines are sample averages prior to 2014 and post 2014.

In this paper, we exploit the variation in housing wealth generated by the diverging trends displayed in Figure 2. However, the decline in oil prices affected other economic conditions in the oil region beyond house prices. Most notably, workers employed in the oil sector (or related to oil extraction or services) faced both increasing unemployment and income risk. To isolate the house price shock from these effects, we focus our analysis exclusively on government sector workers. A central premise is that government workers did not experience *other* shocks in addition to the decline in housing wealth. We revisit this issue at length in Section 3.4.

While our empirical setting is Norway, its institutional features—high homeownership rate, flexible mortgage refinancing, and centralized wage bargaining—make it broadly comparable to other advanced economies with developed mortgage markets. The mechanisms we study, centered on collateral-driven borrowing and the local propagation of housing-wealth shocks, are therefore likely to generalize beyond the Norwegian context.

Figure 2: The development in house price indexes in oil region and rest of Norway.



Note: The left panel displays local house-price indexes for different housing markets across Norway. The oil region markets are marked in red and the rest of Norway in gray. The right panel shows the same indexes, detrended by average year and quarter. Source: Eiendom Norge (Real Estate Norway).

We next turn to our empirical design, which leverages this regional variation to estimate how housing-wealth shocks translate into changes in consumption.

3.2 Reduced-form evidence

We start by exploring the reduced-form effect of living in the oil region on consumption. Specifically, we estimate a difference-in-differences model focusing on consumption as defined in equation (1). We compare the consumption dynamics of government workers living in the oil region in 2014 with government workers living elsewhere, i.e. we estimate:

$$\bar{c}_{i,t} = \sum_{\substack{\tau \in \{s,\ldots,S\}, \\ \tau \neq S_0}} \beta_\tau \, \mathbb{1}(t = \tau) \times \text{oil region}_{i,2014} + \eta_i + \eta_t + \varepsilon_{it}. \tag{2}$$

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 and year-quarter.

In Figure 3 we plot the results. The point estimates of β_τ are interpreted as the difference between the logarithm of total spending between government workers in the oil region relative to government workers in the rest of Norway. As illustrated in the figure, there are no

significant differences in consumption patterns prior to the oil price collapse. However, once oil prices and house prices start to fall, government workers in the oil region exhibit a sharp and statistically significant drop in spending.

3.3 The IV model

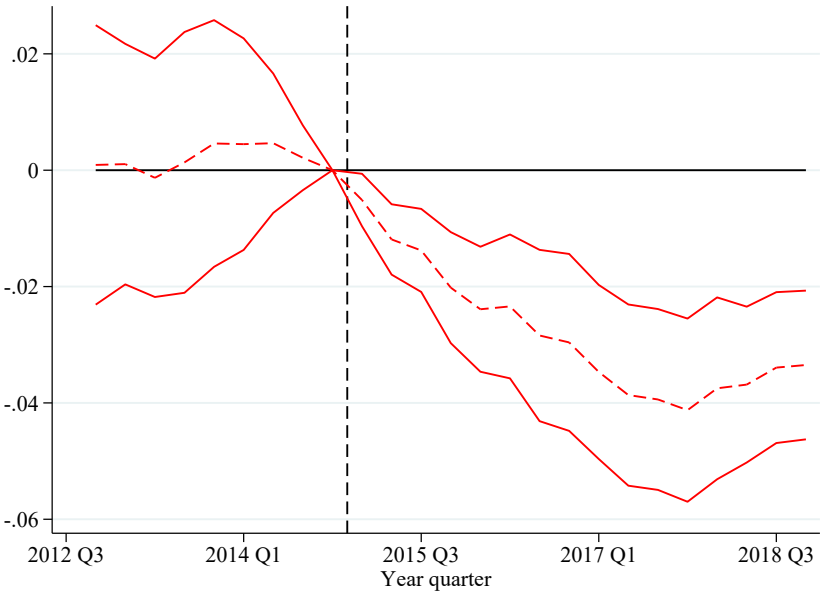
To obtain the MPXs out of housing wealth, we want to estimate

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

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 Norway using an IV approach. The main goal of our IV design is to purge all changes in housing wealth of other correlated effects.

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 (i.e., the data from Figure 2), we compute an estimated value of the home in each quarter from 2012 to 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 $hw_{i,2014+\Delta t}$. Figure 4 plots the implied distribution of household-level home price changes over the 2015–2016 period derived from this procedure. This is the variation in house prices we focus on in our research design. In the remainder of this section, we describe how we use an IV approach to account for the fact that exposure to house price shocks may not be exogenous.

Figure 3: Reduced form diff-in-diff: The spending response of government workers in oil region relative to the non-oil regions.



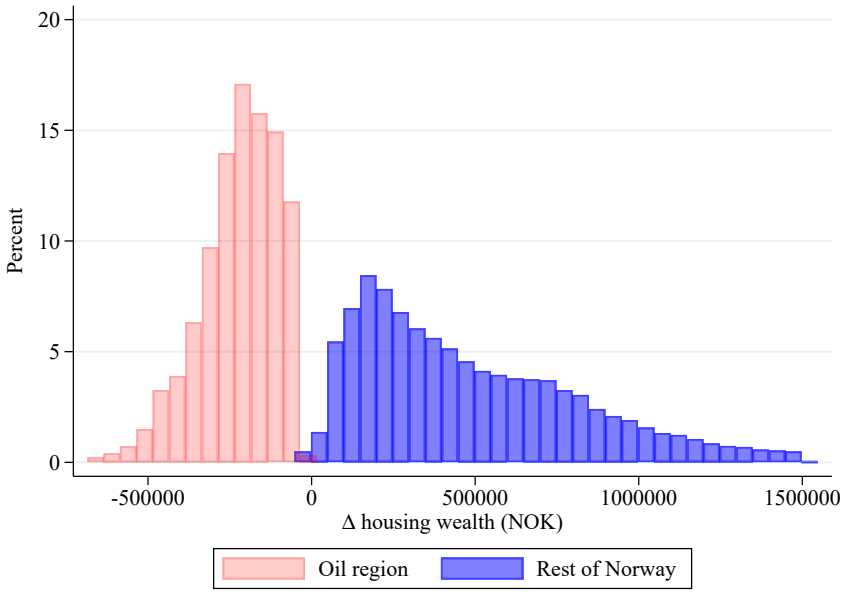
Note: This figure presents the dynamic response of log total spending for government workers in the oil region and government workers in the rest of Norway over time. This is estimated using the difference-in-differences model (2).

First stage In our first-stage regression, we compute the change in housing wealth for household i since 2014 and regress this change on an indicator variable for whether the household resides in the oil-intensive region ℓ (equation (3)). We also control for a set of household characteristics that may be correlated with local home price changes and could independently influence consumption.

$$\Delta hw_{i,2014+\Delta t} \equiv \frac{hw_{i,2014+\Delta t} - hw_{i,2014}}{hw_{i,2014}} = \delta \times \text{oil region}_{i,2014} + \quad \eta_1 X_{i,2014} + \alpha_1 + \varepsilon_{1\,i,\Delta t} \quad (3)$$

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 included as 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 oil region market and is not explained by changes in the characteristics of local government workers.

Figure 4: 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.

The predicted change in housing wealth that we can estimate with (3) is thus the change in the value of households primary's housing due to the location's exposure to the oil price fall. Results from the the first stage IV regression, reported in Table 4, show that our instrument is strongly significant with a first-stage F-test around 300.

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

$$\frac{\bar{c}_{i,2014+\Delta t} - \bar{c}_{i,2014}}{hw_{i,2014}} = \beta_{\Delta t} \widehat{\Delta hw}_{i,2014+\Delta t} + \quad \eta_2 X_{i,2014} + \alpha_2 + \varepsilon_{2\,i,\Delta t} \tag{4}$$

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

Table 4: First stage IV

Dependent variable: $\frac{\Delta hw}{hw_0}$	(1)	(2)
Oil region$_i$	-0.246^{***}	-0.251^{***}
	(0.013)	(0.015)
First Stage F-test (cluster-robust)	344	293
First Stage F-test (Cragg Donald)	38 253	40 220
With additional hh controls	No	Yes
Num. obs.	118 365	118 365

Note: Dependent variable is change in housing wealth between 2014 and 2017. Consumption is smoothed over the last four quarters. Controls used in both stages: labor income, capital income, bank deposits, gross wealth, debt, household size, household head age, LTV, DTI, LTI. All are standardized. All are included as polynomials of power 1–4. Cluster-robust at EN-region level.

variables by four in order to express the coefficient in per-year terms instead of quarterly effects.

We do not observe homeowners' expectations about future home prices. In the specification described in Equation (4), we fix the benchmark year at 2014 and progressively move the end-year forward one quarter at a time. Specifically, we start with the window 2014Q4–2015Q3, then proceed to 2015Q1–2015Q4, and so on, until the end of the sample. In this setup, the change in housing wealth varies gradually over time, accumulating quarter by quarter. As such, we implicitly assume that any observed change in spending reflects the *current* decline in home prices in the oil region relative to the rest of Norway. We refer to these estimates as the "short-run MPX."

However, it is plausible that households anticipate further divergence in regional house prices. To study the implications of such forward-looking behavior, we provide a second set of estimates in which we hold the change in housing wealth fixed at the longest available window in our data. Specifically, we set $\widehat{\Delta hw}_{i,2014+\Delta t} = \widehat{\Delta hw}_{i,2017Q4}$ for all Δt, and estimate what we refer to as the "long-run MPX." This alternative measure facilitates comparison with previous studies that typically rely on "final" or long-run housing wealth shocks.

3.4 Threats to identification

Our exclusion restriction is that residence in the oil region only affects spending through housing wealth. Formally,

$$\mathbb{E}\left[\text{oil region}_{i,2014} \times \varepsilon_{2\,i,\Delta t} \middle| X_{i,2014}\right] = 0 \tag{5}$$

In this section, we discuss two main potential threats to that assumption. First, there may be structural differences between the treated and control households. Second, even if treated and control households are comparable prior to the oil price shock, other confounding shocks could differentially affect the two groups.

3.4.1 Structural differences between treatment and control

A potential concern is that government workers may differ systematically across regions in terms of their consumption patterns or financial positions. For example, distinct local business cycles could place government workers in relatively booming areas on a higher consumption growth trajectory than those in other regions, even if their wages are centrally negotiated. Such differences would bias our estimates if earnings or spending were correlated with local economic conditions independent of housing wealth.

As discussed earlier, Table 2 provides summary statistics for government workers in oil-exposed (treated) and non-oil (control) regions, measured before the oil price collapse. The groups are nearly identical across demographics, income, and liquid wealth. The main difference is that households in oil regions hold somewhat higher housing wealth and debt, leading to slightly higher leverage ratios, but the magnitudes are modest. Consumption levels and composition are also very similar across groups.

Complementing this, Figure 3 shows that, despite the sharp decline in spending after 2015 in the oil region, there are no persistent differences in spending trajectories between oil and non-oil regions during the pre-period. This parallel-trends evidence further supports

the comparability of treatment and control households.

In sum, we find no evidence that government workers in oil-exposed regions differ systematically from their counterparts elsewhere prior to the shock. Observable differences are small, and pre-trends in consumption are parallel, reinforcing the validity of our identification strategy.

3.4.2 Confounding shocks

Next, we discuss the presence of any confounding shocks that could differentially affect government workers following the oil price collapse. A central assumption underlying our empirical strategy is that the estimated consumption response is not due to changes in (expected) labor or capital income among government workers in the oil region. If income expectations deteriorated for government workers in the oil region relative to those in the rest of the country, this would potentially be a confounding factor driving spending. Similarly, if government workers in the oil region were more likely to hold equity in oil-related firms, a decline in the value of these assets following the oil price shock could differentially reduce their wealth and, consequently, their consumption (Chodorow-Reich et al., 2021).

A second potential confounding factor is peer effects. That is, government workers in the oil region may reduce consumption not because of their own financial situation, but because their peers, such as neighboring oil industry workers, cut spending due to job loss or increased precautionary saving.

To assess whether changes in labor or capital income may confound our estimates, we estimate a standard difference-in-differences regression for government worker outcomes, using an oil region indicator as the treatment variable. Specifically, we estimate:

$$y_{i,t} = \beta \, \text{post}_t \times \text{oil region}_i + \eta \, X_{i,t} + \varepsilon_{i,t}, \tag{6}$$

where $y_{i,t}$ is the logarithm of either household i's home value, total consumption, labor

income, and capital income in year t. The vector of controls $X_{i,t}$ contains household mean age and size, and fixed effects for households, housing markets, years in addition to both post_t and oil region_i.

The first two columns in Table 5 replicate the first-stage and the reduced-form consumption evidence from earlier in a panel setting. Government workers in the oil-intensive region experienced a significant relative decline in home values of 25 percentage points. At the same time as house prices decline, government workers in the oil region also reduce consumption by 3.3 percentage points. Importantly, however, labor income did not change differentially across regions. The estimated coefficient in column (3) is small and statistically insignificant, providing strong support for our identifying assumption that government workers in the oil region and the rest of Norway follow similar earnings trajectories. Likewise, we find no evidence that capital income among government workers in the oil region was differentially affected by the 2014 oil price shock.

To complement this, we later compare the spending responses of homeowners and renters within the same regions. Renters face identical local labor-market conditions and exposure to regional demand but do not hold housing assets. If our identification successfully isolates the housing-wealth channel, renters' expenditures should not respond to the oil-price-induced house-price shocks. Consistent with this prediction—and in line with Aladangady (2017), who finds negligible responses for U.S. renters—we find no significant change in renters' spending. These results, reported in Section 4.3, provide a direct falsification test supporting the causal interpretation of our design.

In addition to these tests, we conduct several robustness checks. First, we ensure that we estimate the MPX conditional on a rich set of predetermined observables to ensure that comparisons are made between observationally similar government workers. Second, we repeat our analysis using an alternative sample consisting solely of home-owning retirees. While working-age adults may adjust their consumption in response to changes in expected future earnings, retirees' remaining lifetime income is largely determined by pension entitlements

Table 5: Difference-in-difference model results

	(1) Log(House Val)	(2) Log(Cons)	(3) Log(Labor Inc)	(4) Capital Inc
Post x Oil region	-0.254^{***}	-0.033^{***}	-0.009	-738.000
	(0.015)	(0.010)	(0.006)	(527.000)
R^2	0.989	0.796	0.872	0.322
Household FE	Yes	Yes	Yes	Yes
Region FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Num. obs	820,694	820,549	820,694	820,694

Note: This table presents the results from estimating equation (6). Errors are three-way cluster-robust at the level of household ID, housing market, and year. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

and should be unaffected by local labor market conditions. At the same time, retirees are still exposed to the same housing wealth shocks. If retirees exhibit consumption responses similar to those of working adults, this would mitigate concerns about confounding income effects.

Finally, to shed light on the role of peer effects, we compare neighborhoods where unemployment increased more sharply (and peer effects therefore are likely more salient) with those where unemployment rose less. If peer effects are important, we would expect to observe substantially larger estimated MPXs in the former group and small or non-significant MPXs in the latter group. Results from these robustness exercises are presented in Section 4.3.

4 How does housing wealth affect consumption?

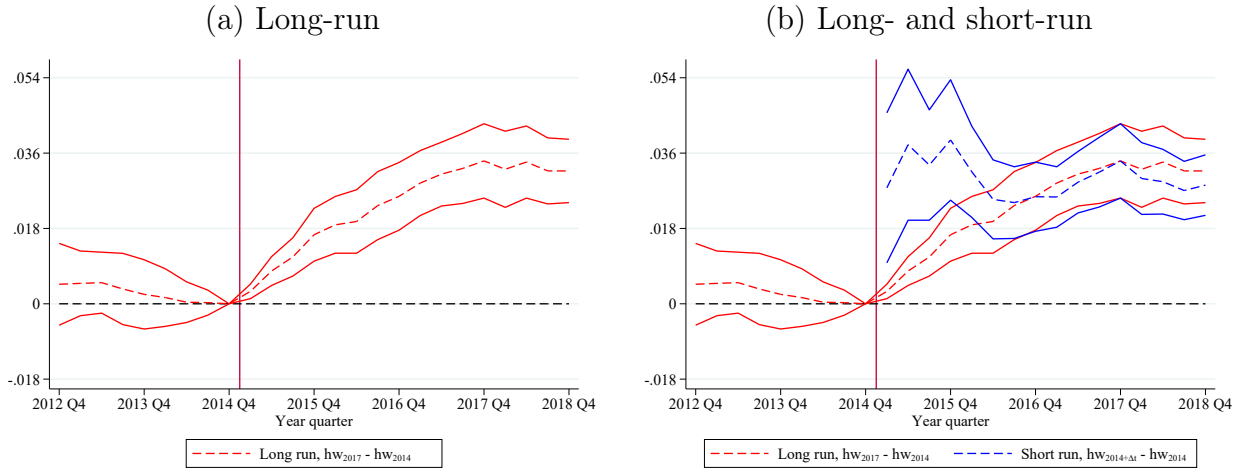
In this section, we present our main results on the size of the spending response to housing wealth changes. We start by reporting the results for the total consumption response, before showing results across disaggregated consumption categories. Motivated by the fact that

only a subset of the consumption categories respond, we define and compute a local MPX which captures the MPX relevant to local producers. We end by showing the results of an extensive range of robustness tests.

4.1 The average consumption responses

In Figure 5 we show the two different estimates of MPXs of total expenditures, i.e., the long-run and short-run MPXs. Each MPX should be interpreted as the annual expenditure effect. The long-run MPX estimates are displayed in red in Figure 5, panel A, while the short-run MPX estimates are represented in blue. Total expenditures fall with decreasing home prices, as expected. The long-run MPX over three years (equal to the current MPX after three years) is 0.036. That is, for each dollar lost in housing wealth, expenditure drops by, on average, 3.6 cents. The magnitude of this change is within the range but at the lower end of earlier estimates. Quantitatively, our estimates are closest to Disney et al. (2010), Guren et al. (2021), and Vestman et al. (2023).

Figure 5: The dynamic marginal propensity to spend out of housing wealth shock



Note: This figure plots the dynamic response of total consumption expenditures for home-owning government workers. Each period corresponds to a four-quarter window beginning at the time shown on the horizontal axis. Short-run MPXs (blue) are based on contemporaneous changes in housing wealth, while long-run MPXs (red) capture cumulative wealth changes over a longer horizon. Shaded areas denote 90% confidence intervals, clustered at the housing-market level.

The dynamic pattern reveals a gradual adjustment: roughly half of the total response materializes within the first year, and nearly the full three-year effect is realized by the second year. This timing suggests that households respond to declining housing wealth both immediately, as liquidity and collateral constraints tighten, and gradually, as they adjust consumption and saving plans. The convergence between short- and long-run MPXs over time reflects this cumulative adjustment process rather than a purely mechanical effect of horizon length.

4.2 The consumption response by consumption categories and local spillovers

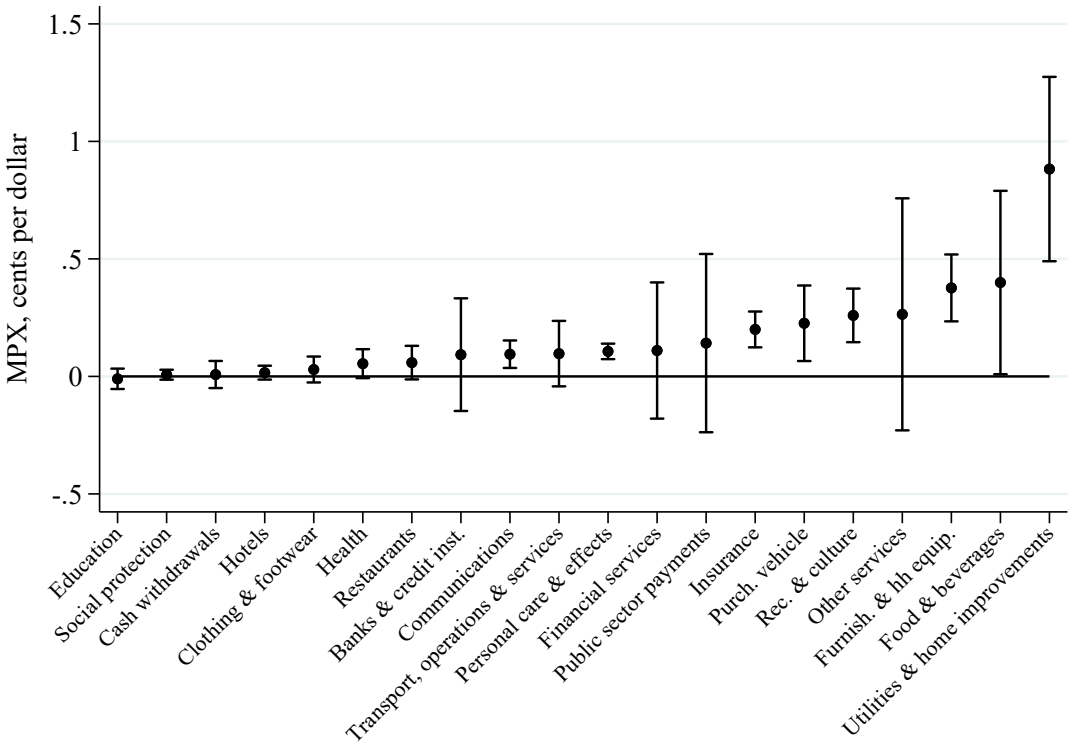
Next, we move on to explore the spending effects across 20 different disaggregated spending categories. We report category-specific three-year MPXs in Figure 6.

As is clear from the figure, the spending response is insignificant for the majority of the spending categories. The types of spending that is affected by changes in housing wealth is typically durable expenditures, particularly in categories such as home improvements, furnishings, and vehicle purchases. However, both food and beverages and recreational and cultural activities also exhibit a significant drop. This implies some adjustment margins for these types of consumption. The remaining categories display either small or statistically insignificant responses. These results highlight an important finding: housing wealth changes transmit to consumption, *but does not lead to a broad-based decline in spending.* One important potential implication of this that the relevant consumption response for local production is lower than our headline estimate.

To explore this, we provide a simple back-of-the-envelope transformation of our headline MPX estimate into a "local-MPX" estimate, which accounts for import leakages.[5] Specif-

[5]A complete evaluation of the macroeconomic consequences of housing wealth shocks requires careful modeling that accounts for general equilibrium effects. This is outside of the scope of this paper. However, it is important for macroprudential policy to understand whether and how house prices affect the local economy. We have documented highly heterogeneous effects across spending categories; since different categories rely on different shares of local inputs, it is not immediately clear how housing wealth shocks impact demand for

Figure 6: 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.

ically, to quantify the first-round impact of changes in housing wealth on demand for local output, we combine three empirical inputs: (i) our estimated marginal propensities to consume out of housing wealth across consumption categories k; (ii) a mapping between consumption categories and producing sectors $B_{k\to n}$; and (iii) sector-level import shares m_n that capture the degree of local production of goods in a sector n.

Consider a small open economy and a change in local housing wealth, ΔHW. The first-round change in spending on consumption category k is

$$\Delta C_k = \text{MPX}_k \, \Delta HW. \tag{7}$$

local production.

Spending in each category translates into demand for sector output according to the bridge matrix $B_{k\to n}$, such that the total demand for industry n generated by category k is

$$\Delta D_{n|k} = B_{k\to n}\,\Delta C_k. \tag{8}$$

Only a fraction of this output is produced locally. Let $1-m_n$ denote the import share of the total use of industry n; the corresponding local-sourcing share is m_n. The locally produced component of the demand from category k is therefore

$$\Delta D^{\text{local}}_{n|k} = m_n \Delta D_{n|k}. \tag{9}$$

Aggregating across industries and categories yields the total first-round effect on local demand:

$$\Delta D^{\text{local,1st}} = \left(\sum_n \sum_k \Delta D^{\text{local}}_{n|k} \right) \tag{10}$$

This framework isolates the immediate, mechanical transmission channel from local housing-wealth shocks to local producers: households cut spending across categories according to their MPCs, and the extent to which this contraction reduces local demand depends on how much of each good or service is produced locally.

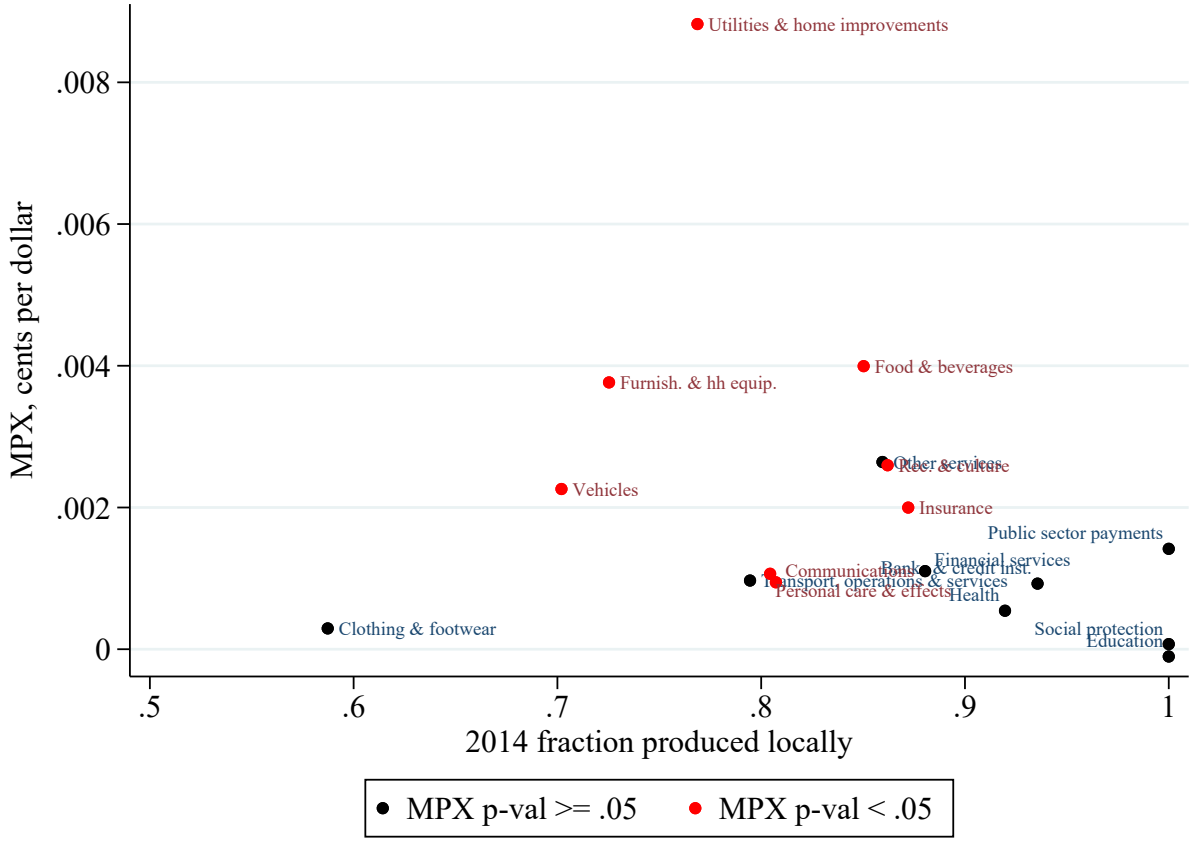
Based on this, we define the *local MPX*, or, the L-MPX, as the demand response relative to the housing wealth change

$$\text{L-MPX} \equiv \frac{\Delta D^{\text{local,1st}}}{\Delta HW} \tag{11}$$

To quantify this, we use input-output tables to proxy m_n as the share of production produced locally. We then compute the mapping $B_{k\to n}$ and use our estimates of MPC_k to compute L-MPX.

In Figure 7, we plot the category-specific MPX against m_n. The figure illustrates that the significantly affected sectors are somewhat intermediate in terms of local sourcing. The most locally sourced consumption categories are associated with insignificant spending re-

Figure 7: MPX and local-sourcing share m_n



Note:

sponses. This illustrates an important point: A large part of the spending response occurs on consumption categories that are at least partially imported. As a result, there is import leakage, which in turn suggests that the local spending response is *lower* than our headline estimate. In fact, in our data, L-MPX is around 80% of our headline estimate.

4.3 Robustness analysis

As discussed in Section 3.4, we conduct a series of robustness exercises to assess whether our estimated housing-wealth effects could reflect selection, compositional changes, or confounding channels such as income shocks, local demand spillovers, or peer effects. These tests

complement our baseline IV estimates and support the exclusion restriction.

We first consider alternative specifications and sample restrictions to verify that our benchmark marginal propensity to spend (MPX) is not sensitive to modeling or subsample choices. Results are reported in Table 6. Column (1) repeats the baseline estimate. Column (2) omits household controls; the MPX increases modestly to 4.9 cents, and the confidence intervals overlap with the baseline. Column (3) excludes Oslo workers to check whether our results are driven by the contemporaneous housing boom in the capital; the estimate remains virtually unchanged. Columns (4)–(5) remove movers to ensure that our results are not driven by moving-related expenditures, while Column (6) shifts the baseline year to 2013 to account for possible anticipation effects. Across all variants, the estimated MPX remains statistically significant and closely centered around our benchmark value.

To further rule out that labor-market conditions drive our findings, we re-estimate the MPX using a sample of retirees. Homeowning retirees are outside the labor market but still exposed to housing-wealth fluctuations, providing a clean test of income-related confounding factors. Due to the specialized age and balance-sheet composition of this sample, we interpret the results cautiously, but they offer a useful complement.

Figure 8 shows that retirees exhibit a positive MPX, albeit smaller in magnitude. Given retirees' lower leverage and higher liquidity, this muted response is consistent with heterogeneity patterns documented in Section 5. Because retirees' labor income is unaffected by local shocks, the similarity in direction of the response strengthens our interpretation that the main results are not driven by income effects.

As previewed in Section 3.4, we also perform a falsification test using renters. Renters face the same local labor-market conditions and regional demand exposure as homeowners but are unaffected by changes in housing wealth. We therefore re-estimate equation (6) using renters' total expenditures.

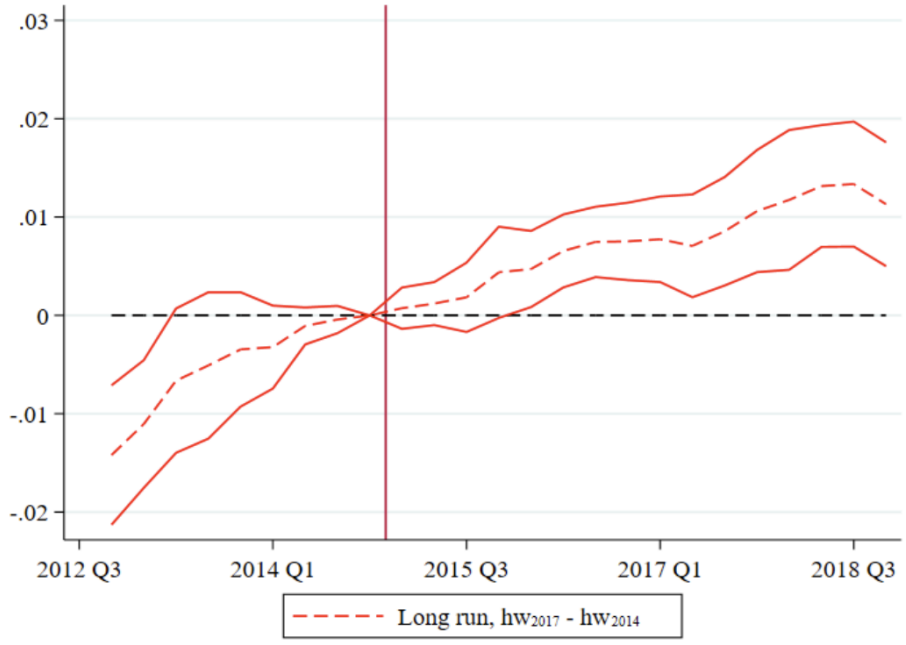
Figure 9 displays the results. Renters show no significant change in expenditures, in sharp contrast to homeowners (Figure 3). Given that renters experience the same local environment

Table 6: Robustness of sample selection

	Benchmark	Without household controls	Excl. Oslo workers	Excl. intra-oil region movers	Excl. all movers	2013 as benchmark year
	(1)	(2)	(3)	(4)	(5)	(6)
MPX	0.034***	0.049***	0.043***	0.041***	0.029***	0.029***
	(0.005)	(0.008)	(0.006)	(0.004)	(0.004)	(0.003)
Num. obs.	117 242	117 242	99 618	116 850	106 376	112 259

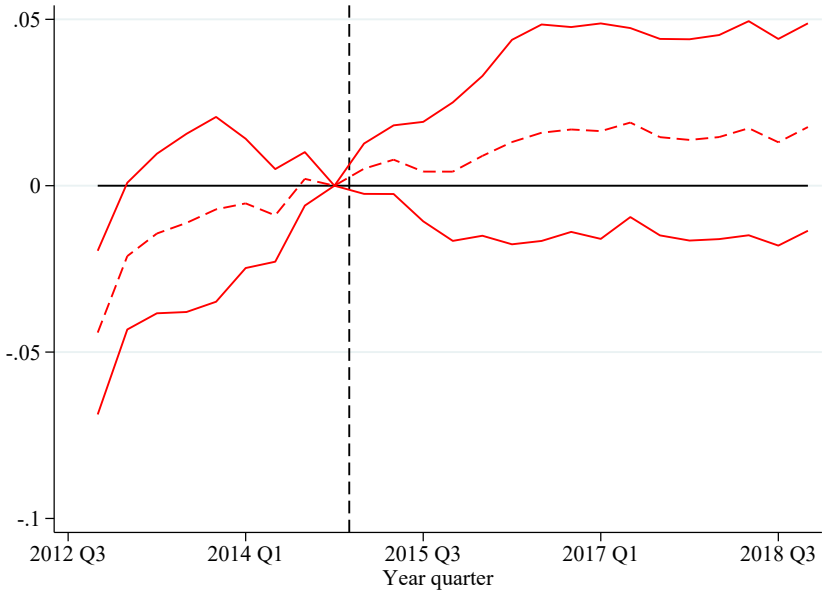
Note: Marginal propensities for different robustness analyses. Column (4) excludes household-year observations with at least one member moving within the oil region during the post period. Column (5) excludes all movers. Column (6) uses 2013 Q4 to compute the housing-wealth change. Observations are fewer because not all homeowners in 2014 were homeowners in 2013.

Figure 8: MPX for retirees



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

Figure 9: Consumption response for renters



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

but no change in housing wealth, this null effect provides strong evidence against local general-equilibrium or demand-spillover interpretations. The absence of a renter response aligns closely with Aladangady (2017), who documents similar findings for the United States.

Finally, we test whether peer effects could bias our estimates. Government workers might cut consumption because their peers—such as oil-sector neighbors—reduce spending due to job losses or heightened uncertainty. To examine this, we proxy neighborhood exposure using the change in local unemployment between 2012–2013 and 2015–2016, and re-estimate the MPX separately for high- and low-exposure areas.

As shown in Table 7, the MPX is only slightly larger in high-unemployment areas, and the confidence intervals overlap substantially, indicating that peer effects are limited.

Taken together, these robustness exercises provide strong evidence that our estimated housing-wealth effects are driven by house price changes and not confounding factors. Retirees, who are insulated from labor-income risk, display similar but smaller responses; renters, who do not hold housing wealth, show zero spending responses; and government workers in low-unemployment areas respond almost as strongly as those in high-unemployment areas. The convergence of these findings supports a causal interpretation in which housing-wealth shocks—rather than income or peer effects—drive the observed consumption adjustments.

Table 7: Investigating peer effects

	High Δunemployment	Low Δunemployment
$\Delta hw/hw_0$	0.039***	0.030***
	(0.005)	(0.006)
Num. observations	114,347	114,347
Num. clusters	105	105

Note: MPX estimates for subsamples defined by above- or below-median local unemployment changes. Standard errors are cluster-robust at the housing-market level.

5 Why do house prices affect consumption?

Our goal in this section is to investigate which dimensions of heterogeneity influence the marginal propensity to spend (MPX) out of housing wealth shocks, thereby uncovering the key mechanisms driving our results. Guided by 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., the impact of changes in home prices on the lifetime wealth of homeowners—has been argued to be relatively small (Buiter, 2010). In theory, housing wealth contributes to consumption only if a housing transaction generates an increase in lifetime resources. The sign of the wealth effect depends on the *expected* future net proceeds from a home sale (Fagereng et al., 2024). However, home sales are often accompanied by purchases—such as when upsizing—which reduce or even eliminate the net value of the transaction. Only in the case of downsizing or moves to less expensive housing markets can homeowners extract net positive value. Conversely, for renters and younger households who expect to purchase housing in the future, rising home prices reduce perceived lifetime wealth. The remaining lifetime also matter for the strength of this channel. For older homeowners, this increase in lifetime resources is smoothed over fewer remaining periods,

potentially resulting in a larger effects if the MPX is driven mainly by wealth considerations.

The collateral effect operates through the relaxation of borrowing constraints for homeowners who use their property as collateral. For credit-constrained households, even modest increases in home prices can significantly raise the nominal borrowing limit. These households—along with those facing liquidity constraints—tend to exhibit higher MPXs in response to transitory income or house price shocks (Parker et al., 2013; Lustig and Van Nieuwerburgh, 2005; Berger et al., 2018). High levels of indebtedness are often correlated with younger age. Thus, if the collateral channel dominates the pure wealth effect, we would expect younger homeowners to exhibit stronger consumption responses than older homeowners.

We explore the relevant mechanisms following two approaches. First, we explore whether balance sheet or demographic heterogeneity are important drivers of the MPX. To the extent that they are, the dimension of heterogeneity can be informative about the underlying mechansim. Second, we explore the extent to which housing wealth changes also leads to balance sheet adjustments by households.

5.1 The role of balance sheet heterogeneity

To disentangle the collateral and wealth channels, we estimate differences in MPXs across household types. Specifically, we construct dummy variables that divide our sample into several groups and interact these with our main explanatory variable. We estimate the following specification:

$$\begin{aligned} \bar{c}^k_{i,2014} - \bar{c}^k_{i,2012} = \sum_{g \in \mathcal{G}} \beta^k_{g,\Delta t} \, \widehat{\Delta hw}_{i,2012+\Delta t} \times z_{g(i),2013} \\ + \eta^k_2 \, X_{i,2012} + \alpha^k_2 + \varepsilon^k_{2\,i,\Delta t}, \end{aligned} \tag{12}$$

Compared with Equation (4), the key difference is the inclusion of the interaction terms involving $z_{g(i),2013}$.[6] In our preferred specification, $z_{g(i),2013}$ denotes time-invariant dummy

[6]Note that, since $z_{g(i),2013}$ is time-invariant, it is absorbed by household fixed effects and should therefore

variables that divide households into two groups based on the value of the heterogeneity variable in 2013—that is, prior to the housing shock.

Our main approach involves constructing dummy variables, assign each household to one of two groups based on these dummies, and estimating separate models for each subsample. We construct four dummy variables. Two of these indicators proxy for credit constraints: the debt-to-income (DTI) ratio and the loan-to-value (LTV) ratio. The remaining two are age (measured as the age of the oldest household member) and liquidity, proxied by total bank deposits.[7] For LTV, households are classified as high or low based on whether their LTV exceeds 90 percent. The threshold for high DTI is set at 4.0, and for age at 46 years. Households are classified as having high (low) liquidity if their bank deposits exceed (fall below) NOK 200,000.[8]

The estimated MPXs for each subsample are presented in the first two rows of Table 8. The coefficient on the difference—i.e., $\beta^k_{g,\Delta t}$ from equation (12)—is reported in the third row. The table reveals substantial variation in the MPX across groups. For households with high loan-to-value (LTV) and high debt-to-income (DTI) ratios, the estimated MPX is approximately twice as large as that for households with low LTV and DTI.

Although the difference is smaller in magnitude and not statistically significant, liquidity displays the opposite pattern. In line with the relevance of credit constraints, the estimated consumption response decreases with increasing liquidity. Finally, the MPX is lower for older households compared to younger ones.

Taken together, the results presented in this section support the interpretation that housing plays an important role as collateral for financing consumption. Households that are likely to be credit constrained exhibit the largest consumption responses, while older households respond less. Notably, even households that are presumably not credit constrained—those

not be included on its own in the regression.

[7]Alternative definitions and cutoffs yield qualitatively similar results. Note that since 2010, Norwegian borrower-based regulations have limited households' maximum DTI and LTV ratios. In 2013, the regulatory thresholds were a DTI below 5.0 and an LTV below 80 percent.

[8]In a separate approach, we divide households into 10 groups based on the heterogeneity variable. These results are consistent with those of our main approach and are reported in Appendix Figure 12.

Table 8: MPX heterogeneity

	LTV	DTI	Age	Liquidity
	(1)	(2)	(3)	(4)
MPX, low	0.030***	0.030***	0.041***	0.038***
	(0.005)	(0.005)	(0.004)	(0.006)
MPX, high	0.057***	0.064***	0.028***	0.029***
	(0.009)	(0.009)	(0.006)	(0.006)
Δ MPX,	0.028***	0.034***	−0.012**	−0.010
high – low	(0.009)	(0.011)	(0.005)	(0.008)
Cutoff:	90%	4.0	45	200 000
Num. obs.	117 242	117 242	117 242	117 242

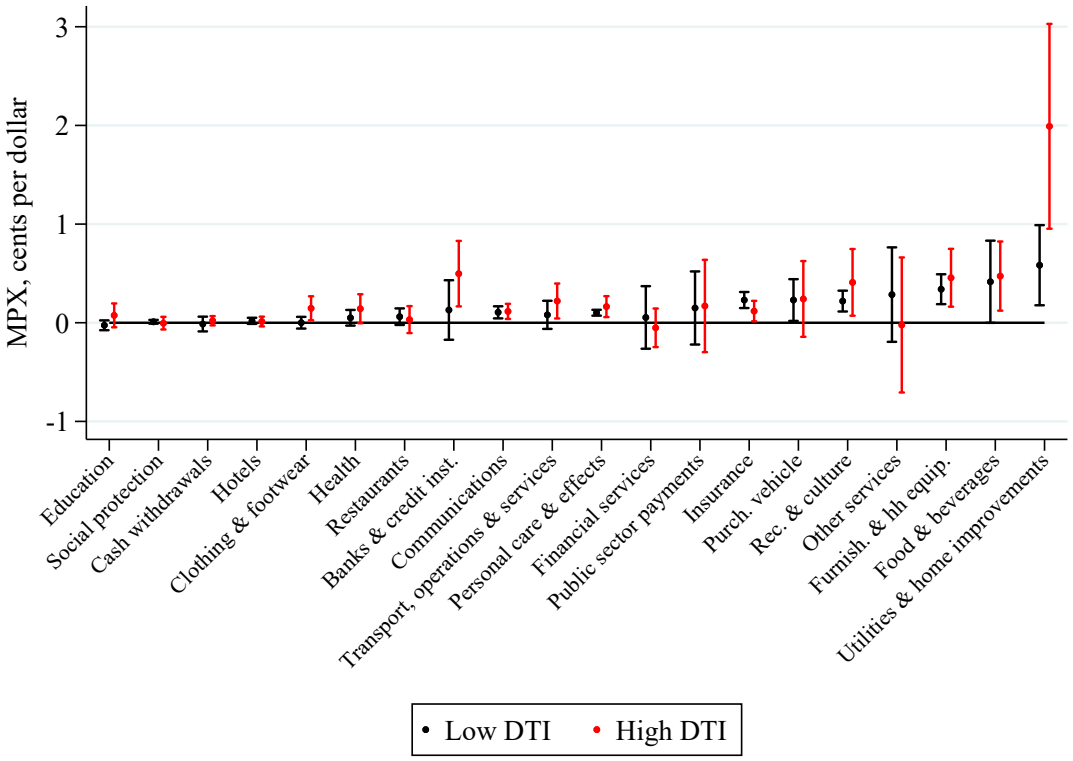
Note: MPX estimates for various subsamples defined by high and low LTV, DTI, age, and liquidity, respectively. We assign high LTV and DTI to households with an LTV above 90% or a DTI above 4. We assign low age to those below 45 years old and low liquidity to those having liquidity below 200 000 NOK.

with low DTI and low LTV—reduce consumption in response to the housing wealth shock. In contrast, retired homeowners display only a weak and statistically insignificant response. On its own, this suggests a limited role for pure wealth effects.

To further support the notion that collateral effects are driving our findings, we proceed by exploring *which* consumption categories drives the heterogeneous responses. In practice, only some consumption categories are financed with collateralized borrowing, and as such the heterogeneity across consumption categories could be informative.

Figure 10 is parallel to Figure 6 but where we estimate the responses for each category of consumption for the two groups with high and low DTI. For most categories, the high-DTI households have a larger estimated MPX but the point estimates are quite similar and the confidence bands suggest that the differences are not statistically significant. For " Utilities and home improvements", however, the difference is very large. This type of spending is perhaps the most likely type of spending to be financed with collateralized lending. These

Figure 10: 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.

results are therefore consistent with the notion that collateral constraints are the important driver of the the MPX.

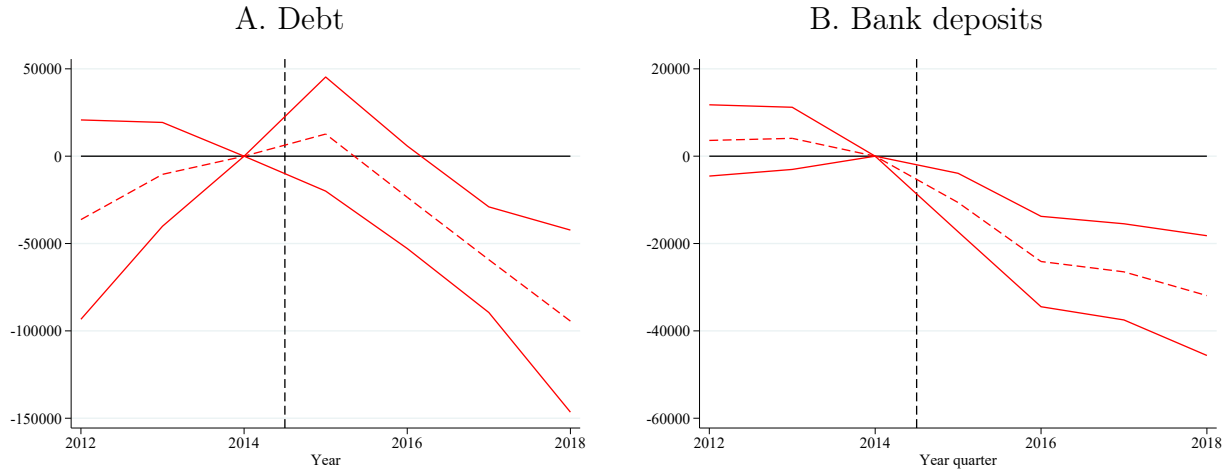
5.2 The effect of house price shock on credit and liquid wealth

We provide additional evidence on the collateral channel by examining the debt response of government workers. If households use their mortgages as a source of credit – through equity withdrawals or home equity lines of credit – we would expect reduced debt accumulation when house prices decline. In addition, when access to credit is constrained, households may resort to using their available liquid assets, leading to a reduction in bank deposits.

To explore this, we estimate a model similar to equation (6), replacing the indicator

post_t with year-specific dummies (excluding 2014), and using total debt and bank deposits as dependent variables. The results are presented in Figure 11. Consistent with our hypothesis, we observe a significant decline in debt growth following 2014 – although the effect is not statistically significant in 2015. Similarly, households appear to reduce their bank deposits relative to the control group following the oil price shock.

Figure 11: 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. (6).

6 Conclusions

This paper quantifies how housing wealth affects household spending using administrative data linked to detailed transaction records for Norwegian households during the 2014–2017 oil price collapse. By comparing government workers in oil-exposed and non-exposed regions, we isolate variation in housing wealth that is plausibly exogenous to income dynamics. We estimate an average marginal propensity to spend out of housing wealth of about 3.6 cents per year over three years. The spending response is highly concentrated in a few durable categories—home improvements, furnishing, and vehicles—indicating that the aggregate effect of housing wealth on consumption is not broad-based. When accounting for local production

shares, the local-MPX is roughly 80% of the headline estimate, implying substantial import leakage and attenuated local demand spillovers.

We show that home equity extraction is the central channel through which housing wealth changes translate into consumption adjustments, with constrained households responding about twice as strongly as unconstrained ones. While Norway's mortgage market is more flexible than most continental systems, the mechanisms we uncover—collateral-driven borrowing and import leakage—are likely to generalize to advanced economies with similar home-equity extraction channels.

Our findings demonstrate that both the composition of spending and the mechanisms of financing matter for how housing shocks propagate through the real economy. The concentration of responses in a few sectors suggests that aggregate housing–consumption linkages may operate primarily through collateralized durables rather than broad consumption smoothing. The evidence also highlights that pure wealth effects remain meaningful, even among unconstrained households, underscoring the need for models that integrate collateral and wealth channels.

Our findings carry direct implications for macroprudential and monetary policy. Because the strongest spending responses arise among highly leveraged households and in a narrow set of collateral-sensitive durables, borrower-based tools such as loan-to-value (LTV) and debt-to-income (DTI) limits can meaningfully reduce local amplification of housing-market shocks. At the same time, our estimates of the local marginal propensity to spend show that a sizable share of the aggregate response leaks to imported or non-local goods, limiting the scope for regional stabilization through housing channels alone. The combination of concentrated collateral effects and substantial leakages underscores that the effectiveness of policy interventions depends jointly on the distribution of household leverage and the structure of domestic production.

Future research should build on these results along three lines. First, given that spending responses are concentrated in a few sectors, the role of the housing market for local demand

during downturns should be revisited using structural models that embed the local-MPX estimates derived here. Second, given the importance of heterogeneity, reconciling estimates across countries and samples through structural approaches that account for sample composition and institutional differences is a key next step. Third, as home equity extraction emerges as the dominant transmission channel, theoretical and quantitative models of housing and consumption should explicitly incorporate borrowing against housing collateral to capture how shocks to house prices propagate to aggregate demand.

References

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

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).

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.

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

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

Auclert, A., M. Rognlie, and L. Straub (2020, January). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Working Paper 26647, National Bureau of Economic Research.

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.

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.

Beraja, M., A. Fuster, E. Hurst, and J. Vavra (2019). Regional heterogeneity and the refinancing channel of monetary policy. *Quarterly Journal of Economics 134*(1), 109–183.

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

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

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

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

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.

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.

Chodorow-Reich, G., P. T. Nenov, and A. Simsek (2021, May). Stock market wealth and the real economy: A local labor market approach. *American Economic Review 111*(5), 1613–57.

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.

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

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

Farhi, E. and I. Werning (2016). Fiscal multipliers: Liquidity traps and currency unions. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, pp. 2417–2492. Elsevier.

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.

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.

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.

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.

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

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.

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

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.

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

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.

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

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

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.

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.

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

Vestman, R., J. Bojeryd, B. Tyrefors, D. Kessel, et al. (2023). The housing wealth effect: Quasi-experimental evidence. Working Paper, Centre for Economic Policy Research.

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).

Our final sample of governments workers are 117 418 households. Below, we lay out how the different sample filtration affect our final number of observations:

Condition	Num. obs.
Initial sample (government workers)	371,464
Self employed income (as share of total income) $< 5\%$	350,301
Strictly more than half in gov worker	165,968
Observed with positive consumption data 2013 and 2014	165,728
Remove top 99.5 of the distribution of total wealth and consumption	146,307
Restrict attention to homeowners	117,418

2 Heterogeneity

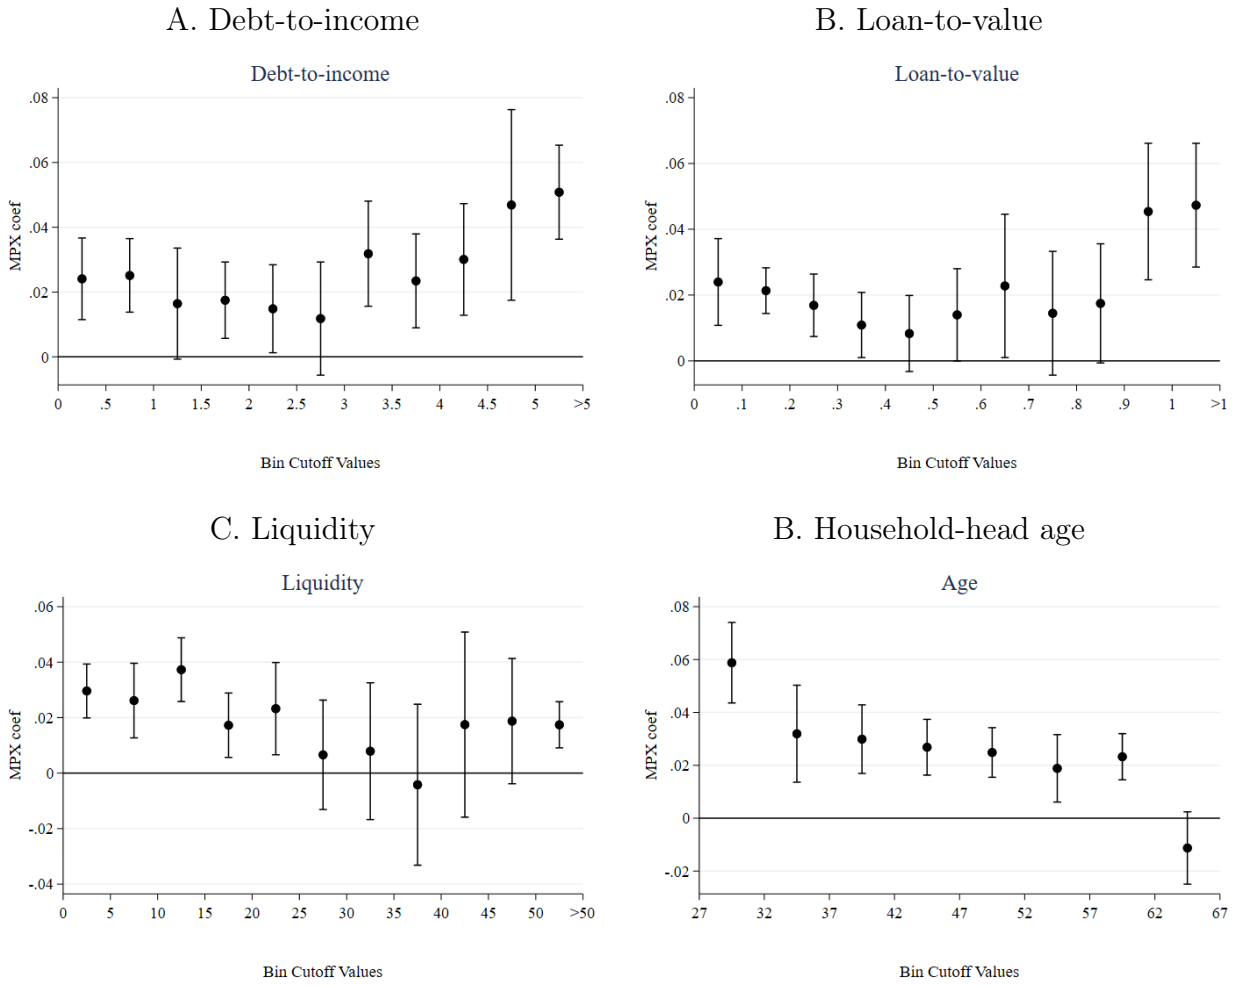
For LTV, we split up the sample in intervals of 10%, i.e., the group with the lowest LTV has a value in the range of 0–0.1, the second-lowest group is in the range of 0.1–0.2, and so forth. DTI divides households in quintiles, age by 6 years, and LTI by units of 2.5.

Table 9: Classification of consumption categories

COICOP group (paper label)	Examples
Food & non-alcoholic beverages	Groceries, staples
Alcohol & tobacco	Beer, wine, cigarettes
Clothing & footwear	Apparel, shoes
Housing services (rent, imputed)	Rent, HOA, fees[a]
Utilities & fuels	Electricity, heating, water
Household furnishings & appliances	Furniture, white goods
Home maintenance & improvements	Materials, contractor services
Health	Out-of-pocket medical, pharmacy
Transport goods (vehicles)	New cars, motorcycles
Transport services	Repairs, public transport, taxis
Communication	Telecom services, equipment
Recreation goods	Electronics, books, sporting goods
Recreation & culture services	Gyms, events, streaming, cinemas
Restaurants & hotels	Dining out, accommodation
Education	Tuition, courses
Personal care	Hairdressing, toiletries
Financial & insurance services	Fees, insurance premia
Misc. goods & services	Small durables, repairs
Postal & delivery	Post, couriers
Other selected durables	Bicycles, tools, baby items

Notes: **(a)** If excluding housing services from consumption (as in some MPC papers), drop this row from all totals.

Figure 12: Heterogeneity in the MPXs



Note: This figure shows the heterogeneity of the total MPX for different household characteristics estimated using (12), while including household controls. Errors are cluster-robust at the level of the housing market and the confidence intervals are of 95%.