

Leverage Regulation and Housing Inequality

Jens Soerlie Kvaerner, Nicola Pavanini, Yushi Peng*

April 2024

Abstract

We estimate an equilibrium model of housing demand and supply. The model allows us to quantify the distributional effects of leverage regulation on mobility and access to high-quality housing. We match the population of households in Norway in 2010-2018, with demographic and financial characteristics, to the universe of housing transactions. Our model features households' dynamic renting and owning choices, investors' housing portfolio rebalancing, and equilibrium pricing across housing products via a market clearing condition. We recover households' willingness to pay for housing quality and moving costs across the income distribution. Our counterfactuals quantify the regressive effects of tighter loan-to-income (LTI) limits, and document how these depend on household preferences and can be offset limiting investors' real estate trading.

JEL: D15; D31; G21; G51; R21; R31

Keywords: Housing demand, housing supply, leverage regulation, housing quality, dynamic discrete choice

*Kvaerner: Tilburg University, jkverner@gmail.com. Pavanini: Tilburg University and CEPR, n.pavanini@tilburguniversity.edu. Peng: Tilburg University, Y.Peng@tilburguniversity.edu. We have benefited from the feedback of Alessandro Gavazza, Victoria Vanasco, Stijn Van Nieuwerburgh, and the seminar and conference participants at Copenhagen Business School, Tilburg, Bank of England, 5th CEPR Endless Summer Conference, PUC Chile, 12th EIEF-UniBo-IGIER Workshop in Industrial Organization, Bank of Italy. Nicola Pavanini gratefully acknowledges financial support from the European Research Council under the Grant Agreement No. 101042069.

1 Introduction

Housing and neighborhood choices impact households' welfare and well-being. Growing up in high-quality neighborhoods positively affects physical and mental health (Ludwig et al., 2012), future college attendance and earnings (Chetty et al., 2016; Chetty and Hendren, 2018b), as well as fertility and marriage patterns (Chetty and Hendren, 2018a). In addition, expected excess returns on real estate wealth are substantially higher than those on financial and pension wealth (Bach et al., 2020). This evidence suggests that access to housing wealth can impact intergenerational mobility and wealth inequality.

Homeownership is only possible for most households via leverage, yet excessive leverage can negatively affect the economy. Mian et al. (2017) provide evidence that in the last 50 years, higher household debt to GDP ratio predicts lower GDP growth and higher unemployment. Mian and Sufi (2009) highlight how excessive lending to subprime borrowers was one of the leading causes of the 2008 financial crisis. To address these negative consequences of the leverage cycle, Geanakoplos (2009) provides theoretical grounds for macro-prudential interventions, suggesting that central banks should regulate leverage with tools such as Loan-To-Value (LTV) and Loan-To-Income (LTI) limits.

While there is evidence of the positive role of macro-prudential regulations on financial stability, as tighter LTI and LTV limits reduce originations of risky mortgages (DeFusco et al., 2020) and improve household debt solvency (van Bekkum et al., 2024), we know little about distributional effects of these policies on household mobility and access to high-quality housing. Uncovering these effects requires separately identifying the role of households' financial constraints from preferences for neighborhood quality in residential choices. This presents two challenges: observing financial frictions requires household-level data on income and wealth, and recovering preferences involves a model of household residential choice.

Our paper addresses these challenges by developing a structural model of housing demand and supply with two novel features. First, it explicitly incorporates households' affordability constraints due to detailed household-level data on income and wealth. Second, it estimates households' heterogeneous preferences for housing, neighborhood quality, and moving costs across the income distribution. We use our model to simulate counterfactual scenarios with tighter LTI limits and document how such policies affect mobility choices and access to high-quality housing across the income distribution. Our model and data allow us to separate the effect of financial constraints from that of households' willingness to pay for quality. We quantify the regressive effects of tighter limits, which depend on households' preferences and the portfolio decisions of real estate investors.

Our framework extends the most recent literature on structural models of residential

choices. Specifically, we generalize the dynamic framework of households' residential location developed by [Bayer et al. \(2016\)](#), introducing five new features. First, we allow two types of real estate investors to demand and supply housing products: financially constrained households, who mostly own and exchange a single housing product and face transaction costs, and financially unconstrained investors, who own portfolios of properties and face no transaction costs. Second, we distinguish between owners and renters among households, allowing renters to become owners and vice versa. Third, all households can stay in their current property, so our analysis not only focuses on those transacting in the housing market during a particular sample period. Fourth, due to household balance sheet data, we can explicitly model households' heterogeneous affordability constraints. Last, we model the equilibrium pricing in housing markets via a market clearing condition that incorporates households' and investors' demand and supply of properties.

Some of our modeling innovations are possible due to administrative data from Norway. The dataset runs from 2010 to 2018 and is geographically restricted to Norway's capital, Oslo. It contains three key components. First, population data is available at the individual-year level for income, debt, wealth, house value, liquid assets, cash, social security payments, demographics, education, and residence location. Family links allow us to merge individuals belonging to the same household, which is necessary to calculate the income and wealth of the household, which are the relevant metrics for mortgage affordability in the data and the model. Second, a transaction-level dataset for the universe of housing transactions in Norway. From 2010 to 2018, we have about 200,000 housing transactions in Oslo. The data includes the unique tax identifier of both buyer and seller for each transaction, allowing us to merge it with the first dataset. The data also contains detailed information on the exact location of the transacted house, its size in square meters, and the transaction price. Third, a set of district-level characteristics from the official statistics published by the municipality of Oslo. We use this data to construct a measure of neighborhood quality.

We define three levels of neighborhood quality and apply the K-means algorithm to allocate neighborhoods to low, medium, and high quality. We use five variables to perform this allocation: education, poverty, crime, health, and happiness. Our quality categorization strongly correlates with house prices and neighborhood residents' income and financial wealth. Furthermore, we document how the life outcomes of individuals between 26 and 35 years old in 2015 are influenced by the quality of the neighborhood they grew up in as children. Growing up in a high-quality neighborhood, relative to a low-quality one, significantly increases the probability of holding a master's degree, having a Ph.D., being a homeowner, and participating in the stock market. Controlling for parents' education as in [Heckman and Landersø \(2022\)](#), homeownership and stock market participation become insignificant

while the educational variables remain significant. Our findings are thus consistent with the hypothesis that parents pass on part of their wealth to children and that neighborhood quality predicts educational attainments across social strata.

To provide descriptive evidence on the importance of leverage regulation, we show the impact of introducing an LTI limit in Norway on households' choice sets. In 2017, Norway set the LTI limit to 5. Before that, there was no strict LTI threshold. We divide housing products across size (small vs large) and quality (low vs high). We show that households in the top 10% of the income distribution, who could afford any house before the LTI limit, can still afford all low- and high-quality small properties after 2017 but lose access to around 10% of low-quality large houses and to about 50% of high-quality large properties. The LTI limit had much larger implications for the choice of less wealthy households. For example, households with median income, who could afford all small properties, lose access to 50% of small low-quality and 70% of small high-quality. In terms of large properties, median-income households could afford 90% of low-quality and 50% of high-quality before 2017, but only 15% of the small low-quality and almost none of the large-high-quality properties post-2017.

We use these descriptive results as motivating evidence for our structural framework, in which we model the housing choices of real estate investors and households. Investors are those with more than one property, which we assume view housing as a financial investment. Households, on the other hand, can be homeowners or renters. They choose whether to stay in their current residence or to move to another one, selected from among a set of housing options, differentiated based on the location and size of the property. Each year, households' housing decisions will be the result of the maximization of their lifetime expected utility, whose preference parameters we will estimate, recovering willingness to pay for house and neighborhood quality across the income distribution. When households move to purchase a house, they incur financial moving costs of both buying the new property and selling the old one. Additionally, when households move as either owners or renters, they incur psychological moving costs. We express these two moving costs as a function of household income and demographic characteristics. The dynamic nature of the model ensures that our estimates of preferences and moving costs reflect a tradeoff between short-term costs and long-term benefits of today's choices.

Households are divided into types, defined by their disposable income, wealth, family size, and age. This categorization allows for heterogeneous choice sets across household types, bound by affordability constraints based on LTI and LTV restrictions, determined respectively by household disposable income and net wealth. If a household moves, its new type reflects the reduction in wealth due to moving costs. Data on homeownership allows us to distinguish between five alternative decisions households can make, which until now have

not been fully captured. First, a homeowner can sell their current property and buy a new one. Second, a homeowner can sell their home and start renting. Third, renters can become homeowners. Fourth, renters can move to another property. Last, a household can choose not to do anything.

Our estimation delivers two sets of results. First, we recover financial and psychological moving costs across the income distribution. Wealthier households face lower costs across both cost dimensions. Older households and larger families face the most considerable psychological moving costs. When recovering the determinants of households' flow utilities that drive the housing choice, we find that households' willingness to pay for housing quality increases with income.

We use our model to simulate a counterfactual scenario with a tighter LTI limit than what was implemented in 2017, setting it to 3 instead of 5. We first document, as expected, that a more stringent LTI limit reduces the share of housing products and the share of high-quality houses differently across the income distribution. The choice sets of the lowest and highest income groups are largely unaffected. For low-income households, this is because the former LTI limit already restricted their choice set to a small fraction of properties. For high-income households, lowering the LTI limit from 5 to 3 is often not enough to restrict their choice set. In contrast to these two groups, households in the middle of the income distribution experience a reduction of up to 18.8% in the share of affordable housing products and up to 26.2% decrease in the share of affordable high-quality properties. Second, we zoom in on households who moved during our sample period, presumably those whose mobility patterns were directly impacted by the implementation of the LTI threshold. For this group, a tighter LTI limit reduces the probability of renters becoming homeowners, equivalent to the probability of buying for first-time buyers, by 11.6% for the lowest income group, and has no effect on the highest income group. In summary, we document the regressive effect of tighter LTI limits, quantifying the reduction in mobility across income and housing quality distributions and the reallocation of households from ownership to rental.

We conduct two additional counterfactuals to propose policies that have the potential to offset these regressive effects of LTI limits. First, we simulate a scenario where we make the real estate investors more inelastic (this can be implemented by, for example, increasing transaction fees on real estate transactions for those with more than one house). This results in an increase in prices especially for properties demanded by high-income households, reflecting their price insensitiveness relative to other household groups, and leads to an increase in household mobility, as investors partially withdraw from trading and households trade more among each other. It also results in a reallocation of mostly low-income households from low- to high-quality neighborhoods. Based on our findings, this happens

because purchasing properties in low-quality districts, which are mostly held by investors, becomes more expensive, so households are better-off renting in higher quality neighborhoods. Second, we simulate a scenario where households' preferences shift closer to those of the top income group. Such a scenario could be implemented through schooling vouchers or public housing assistance programs. This also increases households' mobility and reallocates households from low- to high-quality neighborhoods.

Related Literature. We contribute to three main strands of the literature. First, due to the recent introduction of macro-prudential regulations, empirical evidence on their effects is still scarce and has developed only recently. [Acharya et al. \(2022\)](#) investigate how changes to LTI and LTV limits in Ireland affect mortgage credit and house prices. They find that mortgage credit is reallocated from low- to high-income borrowers and from urban to rural counties, slowing down house price growth. [Peydrò et al. \(2024\)](#) use U.K. mortgage data to show that banks more constrained by a larger exposure to high-LTI mortgages cut credit supply more to low-income borrowers, lowering house price growth. [DeFusco et al. \(2020\)](#) show that the Dodd-Frank Act in the U.S., introducing a rule akin to a tighter LTI limit, has managed to curb originations of risky mortgages substantially. [van Bakkum et al. \(2024\)](#) use Dutch data and a reduction in LTV limits to show that liquidity-constrained households reduce leverage and are less likely to buy a property but have better solvency on their debt. [Han et al. \(2021\)](#) find similar results for Canada. More generally, [Baker \(2018\)](#) documents how heterogeneity in households' consumption elasticity is entirely driven by credit and liquidity, highlighting financial constraints' critical role in addressing household inequality.

Our contribution to this first strand is twofold. First, most of the papers mentioned above focus on how stricter leverage limits affect mortgage outcomes, but none quantifies the effects of these interventions on mobility and access to high-quality housing across the income distribution. Second, the state-of-the-art literature has focused on reduced-form methods that primarily identify local average treatment effects but cannot predict how alternative limits would impact all households across the income and wealth distribution. From a policymaker's perspective, it is crucial to have access to comprehensive predictions on the overall impact of regulatory changes, general equilibrium effects on housing demand and supply, and house prices. This can only be delivered by a structural model of housing demand and supply, which explicitly incorporates LTV and LTI constraints and can predict how household leverage and residential choices would change across counterfactual leverage limits. This is the framework that we develop in this paper.

The second branch of the literature we contribute is on structural equilibrium frameworks to model housing choices. [Bajari et al. \(2013\)](#), [Bayer et al. \(2016\)](#), and [Epple et al. \(2020\)](#) are

examples of dynamic structural housing models.¹ Peng (2021) applies these models to the Chinese housing market, incorporating into the model financing conditions, but only with aggregate data. Almagro and Domínguez-Iino (2024) build a model of residential sorting that quantifies the importance of endogenous location amenities for inequality. An important part of the literature is developing quantitative general equilibrium macroeconomic models that study how financing conditions affect housing choices and equilibrium house prices (Kiyotaki et al., 2011; Sommer et al., 2013; Favilukis et al., 2017). A common limitation of these papers is that they do not consider households’ financing decision. While macro models consider it, they mainly do so from an aggregate perspective, preventing them from analyzing distributional implications with the same degree of heterogeneity as our model.

Last, we contribute to the literature on the effects of social mobility on households’ economic outcomes. Using a randomized housing mobility experiment, recent work documents that households moving to neighborhoods with less poverty experience long-term improvements in physical and mental health (Ludwig et al., 2012), as well as in children’s future college attendance and earnings (Chetty et al., 2016). Chetty and Hendren (2018a) report similar findings, showing that longer exposure to better neighborhoods significantly improves children’s outcomes. Chetty and Hendren (2018b) complement these results with evidence of adulthood increase in income due to exposure to better U.S. counties. Despite these benefits of moving to high-quality neighborhoods, Bayer et al. (2007) find that U.S. households prefer to self-segregate based on race and education. Our framework complements this strand of literature, showing how the combination of leverage regulation and households’ preferences for neighborhood characteristics affect social mobility and the distributional effects of relaxing those regulations.

The rest of the paper is organized as follows. Section 2 describes the data and presents some stylized facts that motivate our model. Section 3 introduces the model, while Section 4 outlines the estimation. Section 5 presents the results, Section 6 shows the counterfactuals.

2 Data and Stylized Facts

In this section, we describe our data sources, define the variables we use, and present a set of summary statistics and stylized facts that motivate our analysis.

¹A comprehensive survey of the literature on structural estimation in urban economics is Holmes and Sieg (2015).

2.1 Data Sources and Variable Definitions

We rely on three data sources. The real estate property and transaction data come from Eiendomsverdi AS (henceforth EV). EV estimates the market value for the Norwegian residential real estate market, both for individual properties and for portfolios of properties.² Our dataset covers the period from 2010 to 2018. The dataset includes all housing transactions in Oslo and several housing attributes. We observe the identity of the buyer and the seller, the date of listing, the transaction price, the number of livable square meters, the number of rooms, and the district. From the municipality of Oslo, we get district-level data on housing stock, rental prices, and district-level characteristics. Rental prices are reported by room number for five aggregate districts. District-level characteristics, which we use to assign a quality score to each district, are available at a more granular level than rental prices.³ We explain all district-level variables and how we use them in Section 2.2.2.

Household-level data are from the Norwegian Tax Registry (NTR) and Statistics Norway (SSB). NTR is responsible for collecting income and wealth taxes in Norway. Employers, banks, and public agencies are obliged by law to submit personal information on income, total assets, and transfers to the NTR before the end of April each year when individuals must submit their tax returns. Individuals are accountable for the information in their tax returns, and the submission of inaccurate information is punishable by law.

We observe each individual’s birth date and the number of children. We merge data on demographics with data on financial information. For each individual, we define income Y as the sum of gross salary and pension plus net capital income and total government transfers. We define net worth, A , as the sum of financial wealth and total assets minus the value assessment of principal residence and debt. We exclude the value assessment of the principal residence because we later add the price of the house to the net worth. We define home ownership as a variable that takes the value of one for all individuals with a positive value assessment of the principal residence. We use the variable $h \in \mathcal{H} = \{0, 1\}$ to separate renters ($h = 0$) from homeowners ($h = 1$).

We distinguish between individuals living alone and individuals with a partner. We obtain the national identity number of the spouse/registered partner from the SSB’s population statistics and use this information to classify an individual into a one-adult household or more than one adult household. We refer to the two household types simply as singles and couples and let G denote the set of demographic variables, including family size, age of the household head, and family type. For tax purposes, the household can allocate wealth in a way that gives the lowest wealth tax. Thus, there are no incentives for tax-motivated asset

²More information is available here: <https://eiendomsverdi.no/>

³All data is freely available here: <https://statistikkbanken.oslo.kommune.no/webview/>

allocation within the household. All households are one-family households.

We calculate the same statistics for both household types. However, for couples, we aggregate total income (Y), net worth (A), and the number of children at the household level. For age and home ownership, we select the maximum in the household. We keep the anonymized identifier of the oldest individual in the household and refer to this individual as the household head. Finally, we require all households to have non-negative financial wealth, debt, and total income and have a household head of at least 18 years of age at the end of each year.

2.1.1 The Housing and Mortgage Market

We now provide the essential institutional details about the Norwegian housing and lending market.⁴ Norway has a high homeownership with a national average of close to 80%. Households, corporations, and the government own these properties. Private landlords, which we refer to as investors (e.g., households with at least two housing units in Oslo), dominate the rental market with a market share of around 80%. Corporations own about 75% of the remaining units, and the rest is owned by the government (Sandlie and Sørvoll, 2017; Stamsø, 2023).

To buy a house in Norway, one typically obtains a pre-qualification letter (“finansieringsbevins”) from a lender that verifies the borrower’s income, performs in-house risk assessments, and ensures compliance with loan-to-value (LTV) and loan-to-income (LTI) requirements. This letter is valid for six months at the time. Following the financial crisis, like many other countries, Norway implemented stricter mortgage regulations, focusing first on loan-to-value (LTV) thresholds. In March 2010, Norway’s Financial Supervisory Authority (FSA) introduced an LTV cap of 90%, which was later reduced to 85% in December 2011 and formalized in 2015.⁵ In December 2016, the FSA introduced an LTI limit equal to five times the gross annual household income. Because the frequency of our data household data is annual, we assume that the LTI limit was introduced in 2017.⁶

The Norwegian mortgage market is dominated by two large banks, DNB and Nordea.⁷ Most mortgages feature variable interest rates with a 20-30-years maturity. Fixed-rate mort-

⁴For more information about the Norwegian housing market, we refer to Brandsaas and Kvaerner (2023) and the references therein.

⁵We refer to Aastveit et al. (2022) for additional details on LTV regulation in Norway.

⁶In practice, many households obtained pre-qualification letters in the months before the introduction of the LTI limit, which means that the LTI constraint did not bind for all households before early 2017.

⁷Their combined market share in the Oslo area was almost 50% in 2020. Nordea became a big mortgage provider after it acquired Danske Bank’s private division. The Norwegian mortgage market is analyzed in a recent report (in Norwegian only) available at https://www.huseierne.no/globalassets/boligfakta/boligfakta-2024/huseierne-det-norske-bankmarkedet-for-boliglan_2024.pdf.

gages are also available, with the most common rate fixation period of 5 and 10 years. Borrowers can compare mortgages online at finansportalen.no. It is a state-financed website run by the Norwegian Consumer Council to provide households with complete information about various financial services, including a mortgage comparison across products and banks. By law, banks and insurance companies must offer their financial products on the portal. The portal has operated since 2008 and has historical mortgage rates since its startup.

For lack of detailed household-level mortgage data, in our model we assume households do not choose which mortgage product to get to finance a property, but rather face a yearly user cost of housing that proxies, among other things, for mortgage payments. Crucially, we assume that the mortgage interest rate component of this user cost does not vary across counterfactuals. We provide some supporting evidence of this assumption in Appendix B.3, where we regress bank-month-mortgage product level interest rates on a set of controls and fixed effects. There we show that 34% of the variation in mortgage rates is driven by the Norwegian Central Bank’s policy rate, and an extra 47% is explained by dummies for mortgage product characteristics (fixed vs variable rates, fixation period, LTV), with bank fixed effects only capturing 1% of the variation. To justify our assumption, we show that, all else equal, the LTI limit introduced in 2017 had no significant effect on mortgage rates.

2.1.2 The Housing Choice Set

Oslo is divided into 18 districts. The 15 largest districts cover approximately 99.5% of the housing stock, so we focus on those. We define the collection of these 15 districts as the set $\mathcal{D} = \{1, 2, \dots, 15\}$. Each district $d \in \mathcal{D}$ is populated with housing units $u \in \mathcal{U}$, distinguished by the number of rooms in the housing unit, $\mathcal{U} = \{1 - 2, 3, 4^+\}$. 1 – 2 includes small housing units with at most two rooms, 3 includes medium units with three rooms, and 4+ includes large units with four or more rooms. We chose this grid to ensure we have multiple transactions for each housing product at each point in time. The Cartesian product $\mathcal{J} = \mathcal{D} \times \mathcal{U} = \{(d, u) | d \in \mathcal{D}, u \in \mathcal{U}\}$ gives a total of $\mathcal{J} = \{1, 2, \dots, 45\}$ housing products. The 45 housing products account for more than 95% of the housing stock in Oslo.⁸

We verify that our discretization scheme of the housing market accounts for a large share of the price variation in the data. Specifically, we decompose the natural logarithm of transaction prices each year into within-product variability and between-product variability. Table 1 presents the results of the variance decomposition. Overall, between 55% and 65% of the total variation in house prices is attributable to between-housing-product variation. We calculate the price $P_{j,t}$ for each product $j \in \mathcal{J}$ in each year t as the average transaction

⁸The excluded fraction of the housing stock (ca. 5%) comes from 3 excluded districts (0.5%) and missing information on the number of rooms (4.5%).

price across all transactions for product j .

Table 1 VARIANCE DECOMPOSITION OF HOUSE PRICES

Year	Total Variance	Between Variance	Within Variance	Between % Total	Within % Total
2010	0.29	0.16	0.13	0.55	0.45
2011	0.28	0.15	0.13	0.54	0.46
2012	0.26	0.14	0.11	0.56	0.44
2013	0.23	0.13	0.10	0.58	0.42
2014	0.23	0.13	0.09	0.59	0.41
2015	0.20	0.12	0.08	0.59	0.41
2016	0.17	0.10	0.07	0.60	0.40
2017	0.17	0.11	0.07	0.63	0.37
2018	0.18	0.12	0.06	0.65	0.35

Notes: This table reports variance decomposition of the natural logarithm of house prices by year. The groups are the 45 housing products constructed as follows: we first define the collection of 15 districts as the set $\mathcal{D} = \{1, 2, \dots, 15\}$. Each district $d \in \mathcal{D}$ is populated with a set of housing units distinguished by the number of rooms in the housing unit, $\mathcal{U} = \{1-2, 3, 4^+\}$. 1-2 includes 1 and 2 rooms apartments and 3+ includes housing units with four or more rooms. The Cartesian product $\mathcal{J} = \mathcal{D} \times \mathcal{U} = \{(d, u) | d \in \mathcal{D}, u \in \mathcal{U}\}$ gives a total of $\mathcal{J} = \{1, 2, \dots, 45\}$ housing products.

2.2 Summary Statistics

2.2.1 The Sample

We construct our sample dynamically. We begin in 2010 and select all households who live in Oslo at the end of the year. For all other years, we consider all households that either lived in Oslo last year or ended up living in Oslo this year. The mobility options are, therefore, moving to Oslo from outside Oslo, staying in Oslo in the same housing product, moving to the same district within Oslo (i.e., only changing house size), moving to another district within Oslo, or leaving Oslo. We refer to the latter as the outside option. Details about the sample construction are in Appendix B. We have approximately 580,000 unique households from 2010 to 2018 and roughly 3.2 million observations. Table 2 presents descriptive statistics for 2015.

2.2.2 Identifying Neighborhood Quality

We use three levels of neighborhood quality and apply K-means to generate three quality clusters. Districts of the same quality have comparable scores on five indicators. These indicators are GPA of primary school, which is the sole criterion for admission to upper

Table 2 DESCRIPTIVE STATISTICS

Panel A: Demographic and Financial Data					
	Homeowners				
	Mean	Std Dev	10th	50th	90th
Fraction couples	0.73	0.44	0.00	1.00	1.00
Age	46	16	28	42	69
Number of children	1.2	1.2	0.0	1.0	3.0
Total income	699	1,553	262	540	1,075
Gross wealth	5,010	20,767	2,529	3,590	7,248
Debt	1,575	2,196	14	1,277	3,362
	Renters				
	Mean	Std Dev	10th	50th	90th
Fraction couples	0.43	0.50	0.00	0.00	1.00
Age	37	15	23	32	59
Number of children	0.73	1.24	0.00	0.00	2.00
Total income	329	693	36	279	594
Gross wealth	301	5,205	0	39	448
Debt	282	1,025	0	50	560
	Homeowners		Renters		
	N of Obs	Share	N of Obs	Share	
Stayers	167,146	0.83	101,476	0.60	
Leaving Oslo	7,031	0.03	8,669	0.05	
Entering Oslo	7,189	0.04	26,368	0.16	
Move within district in Oslo	8,037	0.04	10,519	0.06	
Move between district in Oslo	13,181	0.07	21,453	0.13	
Total	202,584	1.00	168,485	1.00	

Notes: This table reports descriptive statistics of our sample in 2015. Panel A reports demographic and financial data. Panel B reports mobility statistics. We report the descriptive statistics for homeowners and renters separately. All financial variables are reported in NOK thousands.

secondary school, the number of reports to the child welfare service per capita,⁹ criminal offenses by individuals between 0 and 17 years of age per capita, and answers to two survey questions. The first survey question is a self-assessment of health.¹⁰ The second survey question is the score on a happiness index.¹¹ We use the scores at the end of 2015 because it is the only year we observe scores on all five quality indicators.

Table 3 presents the results from a cross-sectional regression of neighborhood characteristics on a constant and two dummy variables for neighborhood quality. Each regression includes 15 data points. The low (high) quality dummy variable takes the value of one if a district belongs to the low (high) quality neighborhood. The constant serves as a reference point. We refer to it as the baseline.

Table 3 HOUSING QUALITY REGRESSIONS

	Outcome variable				
	Price	Income	Financial Wealth	Age	N of Children
Baseline	5,191.56*** (275.76)	420.06*** (14.18)	110.60*** (14.58)	39.43*** (1.56)	0.74*** (0.15)
Low Quality	-1,670.54*** (356.00)	-42.20** (18.30)	-39.66* (18.83)	4.26* (2.02)	0.53** (0.20)
High Quality	1,625.81*** (369.97)	67.62*** (19.02)	72.09*** (19.57)	5.97** (2.10)	0.45** (0.21)
Districts	15	15	15	15	15
R ²	0.87	0.74	0.73	0.31	0.30

Notes: This table reports the results from regressing household and neighborhood characteristics on a constant and two dummy variables for neighborhood quality. The “Low Quality” (“High Quality”) dummy variable takes the value of one if the district belongs to the low (high) quality area. The “Baseline” represents the average score of a middle-quality district. We include the following characteristics as dependent variables: Price for a 3-room apartment, median income and wealth, age and number of children.

The first column reveals large price differences for a 3-room apartment between neighborhoods. The high-quality neighborhood is about 30% more expensive than the baseline,

⁹The total number of reports by district by year is divided by district population and then multiplied by 1,000. This normalization is applied to make all the variables in comparable units.

¹⁰The question is: “How satisfied or dissatisfied are you with your health?” The options are: “Very dissatisfied”, “Slightly dissatisfied”, “Neither satisfied nor dissatisfied”, “Slightly satisfied”, or “Very satisfied”. We count those who answer “Slightly satisfied”, “Very satisfied” as a proportion of all those who responded.

¹¹The question is: “How satisfied are you with your local environment?”. The options are: “Very dissatisfied”, “Slightly dissatisfied”, “Neither satisfied nor dissatisfied”, “Slightly satisfied”, or “Very satisfied”. We count those who answer “Slightly satisfied”, “Very satisfied” as a proportion of all those who responded.

while the low-quality neighborhood is about 30% cheaper. This shows that households’ willingness to pay for a 3-room apartment is increasing in our measure of neighborhood quality. The second and third columns show that the average household income and financial wealth increase monotonically with neighborhood quality. In contrast, the age of the household and the number of children do not.

2.3 Neighborhood Quality and Life Outcomes

With a measure of neighborhood quality, we can calculate the correlation between the quality of the neighborhood people grow up in and their outcomes later in life, in line with [Chetty and Hendren \(2018a\)](#). Following [Heckman and Landersø \(2022\)](#), we also control for parents’ education. We restrict the analysis to 2015 and include ten cohorts between 26 and 35 years old. In 1990, these people were ten years old or younger. We run two sets of cross-sectional regressions, with dummies for whether an individual i in 1990 was resident in a low or high-quality neighborhood, and use the middle-quality neighborhoods as the reference group (i.e., $D_{i,1990} = 0$) in both specifications:

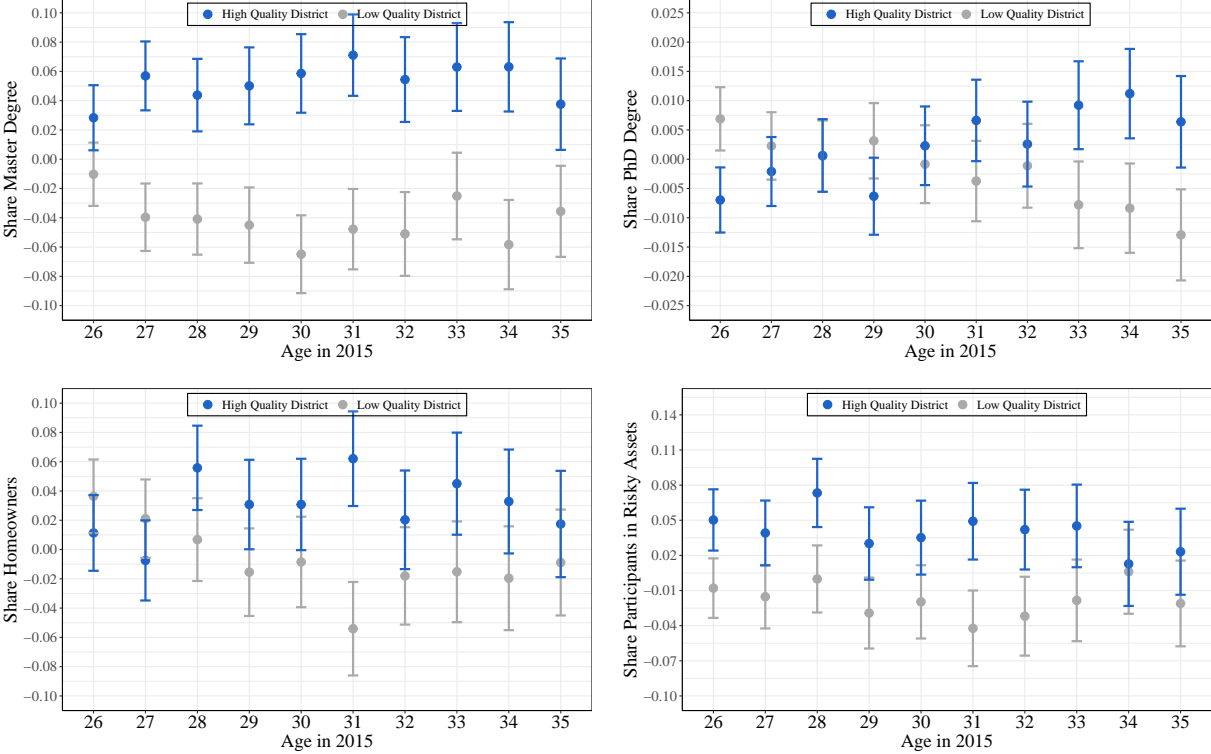
$$I_{i,2015} = \sum_{j=26}^{35} \gamma_j \mathbb{1}\{age_{i,2015} = j\} + \sum_{j=26}^{35} \eta_j \mathbb{1}\{age_{i,2015} = j\} \times D_{i,1990} + \delta E_{i,1990} + \epsilon_{i,2015}. \quad (1)$$

Here $E_{i,1990}$ are dummies for each parent’s years of education in 1990, and $I_{i,2015}$ is an indicator variable that measures life outcomes in 2015. We include four life outcomes: having a master’s degree, a Ph.D. degree, being a homeowner, and participating in the stock market. In our sample, 26% of individuals have at least a master’s degree, 1% have a Ph.D., 55% are homeowners, and 58% are stock market participants. We use the middle-quality neighborhood to identify γ_j . It measures the proportion of individuals at a given age in a middle-quality neighborhood with a dependent variable of one. The coefficients of interest are η_j ’s. These coefficients measure the difference in the proportion of individuals with a dependent variable of one at a given age in 2015 who grew up in either a high or a low-quality neighborhood relative to a middle-quality neighborhood.

Figure 6 presents the results. The key takeaway is that, for most cohorts, growing up in a high (low) quality neighborhood is associated with significantly higher (lower) educational achievements. The corresponding results for measures of wealth, such as homeownership and participation in risky assets, are statistically weaker. In Appendix A, we report the same figures based on regressions without controlling for parents’ education. With this specification, most outcomes are highly statistically significant.

To sum up, our regressions reveal that parents’ education explains most of the difference

in homeownership and stock market participation between high and low-quality districts, consistent with parents' passing on their wealth to children. In contrast, neighborhood quality remains a strong predictor of educational attainment after controlling for parents' education. Taken together, our results indicate that neighborhood quality plays a pivotal role in social mobility.




Notes: The figure plots the regression coefficients $\hat{\eta}_j$ from equation (1) with confidence intervals.

Figure 1 LIFE OUTCOMES WHEN GROWING UP IN LOW VS HIGH-QUALITY NEIGHBORHOOD

2.4 The Effects of Mortgage Regulation on Prices and Choice Sets

Our last preliminary analysis shows the effect of the loan-to-income (LTI) cap on house prices and individuals' choice sets. The event we study is the introduction of an LTI cap in Norway in 2017. Figure 2 shows that the LTI regulation successfully reduced house price growth. The left plot shows the relative price growth. The right plot shows the 6-month rolling mean of average monthly transaction prices. Both plots include all transactions as well as four housing categories. These categories are based on two size groups (1 and 2-room apartments and 3 rooms or larger) and two district qualities (high and low).

We next analyze how the LTI regulation affected individuals' choice sets. We report the relative choice sets available for two sets of households. The first is the median household, and



Notes: These figures plot the cumulative price growth (left panel) and the price level (right panel) across time of housing products of different sizes and quality.


Figure 2 EFFECT OF LTI REGULATION ON HOUSE PRICES OVER TIME

the second is the top 10% households measured by their location in the income distribution. We define the relative choice as the affordable fraction of housing transactions within a month, given the household’s resources and the prevailing mortgage regulations. We focus on the period from January 2016 to December 2018. We use the four housing categories of Figure 2. As explained in section 2.1.1, the LTI cap applied to everyone from June 2017 onwards. Before then, the only requirement for a mortgage was a maximum LTV of 85%. We use the equity and income of the median and top 10% of households at the year’s end to calculate these choice sets. Figure 3 presents the results.

3 Model

The evidence in Table 3 suggests that high-income and high-wealth households are more likely to live in high-quality districts. This could be consistent with heterogeneity across households’ income and wealth distribution in either preferences for neighborhood quality, or in affordability constraints, or both. These patterns also raise the question of how heterogeneity in households’ preferences and affordability constraints drive mobility patterns and house prices. Addressing these questions is important for at least three reasons. First, because the neighborhood where individuals grow up during childhood is a strong predictor of their future sociodemographic outcomes. Second, because housing wealth is the main source of wealth for the median household. Last, because it helps shed light on the distributional effects of leverage regulation. To answer these questions, we develop an equilibrium model of housing demand with mortgage affordability constraints, housing supply, and equilibrium prices.

We consider $i = 1, \dots, \mathcal{N}$ potential buyers and sellers of residential real estate. The



Notes: These figures plot the share of housing products available in households' choice sets across time, before and after the LTI limit applied to everyone as of June 2017. The top two figures refer to households in the 10% of the income distribution, while the bottom two figures refer to the median households in terms of income. The two left figures refer to small properties (1-2 rooms), while the two right figures refer to large properties (3 or more rooms). The solid vertical lines within each figure refer to high-quality properties, and the dashed vertical lines refer to low-quality properties.

Figure 3 EFFECT OF LTI REGULATION ON HOUSING CHOICE SETS

market participants are either households or investors and make a housing decision every year $t = 1, \dots, T$.¹² Households own at most one property in Oslo, and can own other properties (such as a holiday home) outside Oslo. Investors own more than one residential property at any point during the sample period. We group households based on their type $\tau := \tau(Z_{i,t})$, where the variable $Z_{i,t}$ includes households' total income ($Y_{i,t}$), net worth ($A_{i,t}$), and demographics ($G_{i,t}$) such as family size and age of the household head.¹³

¹²We use the period 2011 to 2018 for the estimation and 2010 to determine the initial housing allocation.

¹³To have a finite number of household types, we discretize the variables that define types as follows. We divide income into five groups (all numbers are in NOK, with 1 USD corresponding to about 8 NOK in 2015): <200k, 200k-400k, 400k-600k, 600k-800k, 800k-2,000k, net worth into five groups: <100k, 100k-1,000k, 1,000k-2,500k, 2,500k-5,000k, 5,000k-15,000k, family size into two groups: one-adult household ("singles") or more than one adult household ("couples"), and age into three groups: 18-34, 35-49, and above 50.

3.1 Households

In every period t , a household of type τ decides whether to rent or buy ($\mathcal{H} = \{0, 1\}$) and a housing product $j \in J_{\tau,t}$. More specifically, households in the model can make any of the following choices every period. Renters can choose to stay in their current house, move to another rented property, or buy a house of type j and become a homeowner. Similarly, homeowners can stay in their current property, sell their house to buy another one, or sell their property to become a renter in another house. The housing options differ by location (15 districts) and size (1-2, 3, or 4+ bedrooms). Households deciding to leave Oslo chose the outside option ($j = 0$). Summing up, the choice set of housing options is $\mathcal{J}_{\tau,t} = \{0, \dots, J_{\tau,t}\}$. We denote household's i decision by $d_{i,t} = \{j, h\} \in \mathcal{D}_{\tau,t} = \mathcal{J}_{\tau,t} \times \mathcal{H}$. Moving occurs when a household changes its current housing option, i.e., $d_{i,t} \neq d_{i,t-1}$.¹⁴

We denote a household's new type after the housing decision as $\bar{\tau} := \tau(\bar{Z}_{i,t})$, which allows moving and transaction costs to impact household wealth. To bound households' housing choice set, we use data on income and wealth for each type τ in every period t , together with the average house prices $P_{j,t}$ across locations and sizes, and the actual Loan-to-Income (LTI_t) and Loan-to-Value (LTV_t) thresholds. Suppose a household chooses to purchase a property. In that case, the affordable options must have a price satisfying both the LTI and LTV constraints:¹⁵

$$\mathcal{J}_{\tau,t} = \left\{ j \mid P_{j,t} \leq \min \left(\frac{A_{i,t}}{1 - LTV_t}, \frac{Y_{i,t} LTI_t}{LTV_t} \right) \right\}. \quad (2)$$

We assume that households can afford to rent any housing product.¹⁶ Given a set of affordable options, each household makes a sequence of housing decisions $\{d_{i,r}\}_{r=t}^T$ to maximize lifetime expected utility:

$$\max_{\{d_{i,r} \in \mathcal{D}_{\tau,r}\}} \mathbb{E} \left[\sum_{r=t}^T \beta^{r-t} u(P_{j,r}, X_{j,r}, \xi_{j,h,r}, \chi_r, Z_{i,r}, d_{i,r}, d_{i,r-1}, \varepsilon_{i,j,h,r}) \mid P_{j,t}, X_{j,t}, \xi_{j,h,t}, \chi_t, Z_{i,t}, d_{i,t}, \varepsilon_{i,j,h,t} \right]. \quad (3)$$

Here β is the discount factor, $P_{j,t}$ is the average price, $X_{j,t}$ includes observed house and

¹⁴For around 7% of the transactions, households move to a property of the same size and in the same district as their previous one. We discard these transactions and classify these households as stayers. These are moves that would not occur in the model as no one would pay moving and transaction costs without receiving a benefit.

¹⁵While LTI constraints are straightforward to compute based on our detailed data on household income, LTV constraints are harder to measure. The reason is that households may receive wealth from relatives to purchase a property (Benetton et al., 2022), which we do not observe.

¹⁶We motivate this assumption based on our data, where we observe households of all types renting all types of property.

neighborhood characteristics (house size and neighborhood quality), $\xi_{j,h,t}$ are unobserved house and neighborhood characteristics, χ_t are the LTI_t and LTV_t thresholds, $Z_{i,t}$ are household demographics (income, wealth, age, size), and $\varepsilon_{i,j,h,t}$ is the latent demand of household i for housing product j with the ownership status h , distributed as Type 1 Extreme Value.

We assume the household’s problem follows a Markovian structure and has an infinite horizon. This assumption makes it possible to express the present value of lifetime expected utility as the sum of current utility and the present discounted value of future utility:

$$V(Z_{i,t}, \chi_t, \varepsilon_{i,t}) = \max_{d_{i,t} \in \mathcal{D}_{\tau,t}} \{u(P_{j,t}, X_{j,t}, \xi_{j,h,t}, \chi_t, Z_{i,t}, d_{i,t}) - \mathbb{1}_{\{d_{i,t} \neq d_{i,t-1}\}} F(Z_{i,t}, d_{i,t}, d_{i,t-1}) + \varepsilon_{i,t} + \beta \mathbb{E} [V(Z_{i,t+1}, \chi_{t+1}, \varepsilon_{i,t+1}) | \chi_t, Z_{i,t}, \varepsilon_{i,t}, d_{i,t}]\}, \quad (4)$$

where $F(Z_{i,t}, d_{i,t}, d_{i,t-1})$ represents moving costs. Following [Rust \(1987\)](#), we assume additive separability between per-period utility, moving costs, and the unobserved state variable, as well as conditional independence between the Markovian transition processes of the observed and unobserved state variables.

3.2 Real Estate Investors

Real estate investors own all properties not owned by households. The investors hold portfolios of at least two residential properties and face no affordability constraints or transaction costs. They trade freely each period. We index variables and parameters that apply to real estate investors by s .

Two main differences exist between the modeling of households’ moving decisions and investors’ portfolio choices. First, while a household’s moving decision is a discrete choice among mutually exclusive alternatives, investors solve a (housing) portfolio problem. Second, for households, it is essential to model the dynamic dimension of the housing choice following [Bayer et al. \(2016\)](#), as they trade-off the sunk cost of moving versus their valuations of housing products in all future periods. Because a similar dynamic model with non-mutually exclusive alternatives would be challenging to solve, we model investors’ decisions as a one-period portfolio choice problem inspired by the demand-based asset pricing literature developed by [Kojien and Yogo \(2019\)](#).¹⁷

¹⁷The modeling choice can be justified with, for example, a mean-variance optimizing investor and i.i.d housing returns.

With a one-period objective function, we sum up the holdings of investors to one aggregate investor portfolio. Together, the investor and the household sector own all properties in the data and the model.

3.3 Econometric Model

We define $t_{0,i}$ as the first period we observe a household in our sample and T_i as the total number of periods we observe household i . Let $v_{j,h,t}^\tau = \bar{V}(Z_{i,t}, \chi_t, d_{i,t} = \{j, h\})$ denote the choice-specific expected value function for a household with characteristics $Z_{i,t}$ and decision $d_{i,t} = \{j, h\}$. To simplify the notation, we use τ to denote the household type $\tau(Z_{i,t})$. Let $u_{j,h,t}^\tau$ stand for the deterministic component of flow utility for households of type τ . If a household moves, its new type will be $\bar{\tau} := \tau(\bar{Z}_{i,t})$, reflecting the reduction in wealth due to moving costs. As standard in discrete choice models, normalization is necessary to identify the vector of lifetime utilities $v_{j,h,t}^\tau$. In our case, we estimate a normalized lifetime utility $\tilde{v}_{j,h,t}^\tau = v_{j,h,t}^\tau - m_t^\tau$, where m_t^τ is a normalizing constant that reflects the average lifetime utility of household type τ in time t . In the next section, we will discuss how to estimate m_t^τ in detail.

3.3.1 Households

We start by characterizing the households' decisions. Household i who is considering moving and becoming type $\bar{\tau}$, chooses option j if $\tilde{v}_{j,h,t}^{\bar{\tau}} + \varepsilon_{i,j,h,t} > \tilde{v}_{k,h,t}^{\bar{\tau}} + \varepsilon_{i,k,h,t} \forall k \neq j$. Conditional on moving to an inside option (i.e., for $j \neq 0$), the probability of a household of type $\bar{\tau}$ choosing housing product j with ownership h in period t is:

$$\Pr_{j,h,t}^{\bar{\tau}} = \frac{\exp(\tilde{v}_{j,h,t}^{\bar{\tau}})}{\sum_{l=0}^1 \sum_{k=1}^{J_{\bar{\tau}t}} \exp(\tilde{v}_{k,l,t}^{\bar{\tau}})}. \quad (5)$$

Let $t_{1,i}$ denote the period household i decides where to move (conditional on moving to an inside option). We denote the household's likelihood contribution for this choice by $L_i^{\text{prod}}(\tilde{v})$, where \tilde{v} is the vector of all values of $\tilde{v}_{j,h,t}^{\bar{\tau}}$, which is given by:

$$L_i^{\text{prod}}(\tilde{v}) = \prod_{h=0}^1 \prod_{j=1}^{J_{\bar{\tau}t}} \left(\Pr_{j,h,t_{1,i}}^{\bar{\tau}} \right)^{\mathbb{1}_{[d_{i,t_{1,i}} = \{j,h\}]}}. \quad (6)$$

The probability that a household chooses the outside option in period t , conditional on moving, is:

$$\Pr_{0,t}^{\bar{\tau}} = \frac{\exp(\tilde{v}_{0,t}^{\bar{\tau}})}{\exp(\tilde{v}_{0,t}^{\bar{\tau}}) + \sum_{l=0}^1 \sum_{k=0}^{J_{\bar{\tau}t}} \exp(\tilde{v}_{k,l,t}^{\bar{\tau}})}. \quad (7)$$

Let $t_{2,i}$ denote the period household i considers the outside option (conditional on moving). The likelihood, denoted by L_i^{out} , is:

$$L_i^{\text{out}}(\tilde{v}) = \Pr_{0,t_2,i}^{\bar{\tau}} \mathbb{1}_{[j=0]} (1 - \Pr_{0,t_2,i}^{\bar{\tau}})^{\mathbb{1}_{[j \in \{1, \dots, J_{\bar{\tau}t}\}]}}. \quad (8)$$

In any given period, a household will not move from its current housing product if the indirect utility of staying exceeds the utility value of the best alternative. Recalling that if a household of type τ moves, we denote its new type by $\bar{\tau}$. A household who is currently in house type j with property ownership h will choose to stay if:

$$v_{j,h,t}^{\tau} + \varepsilon_{i,j,h,t} > \max_{\{l,k\} \in \mathcal{D}_{\tau,t}} [v_{k,l,t}^{\bar{\tau}} + \varepsilon_{i,k,l,t}] - \text{PMC}_{i,t}^{\bar{\tau}}, \quad (9)$$

where $\text{PMC}_{i,t}^{\bar{\tau}} = \bar{Z}'_{i,t} \gamma_{\text{pmc}}$ represents the psychological moving cost for a household of type $\bar{\tau}$. These are any costs households incur on top of the monetary cost of moving, which we introduce below. The psychological moving costs depend on the demographics that define the household type. By the definition of the normalized choice-specific value functions, we have $\tilde{v}_{j,h,t}^{\tau}$, where $\tilde{v}_{j,h,t}^{\tau} = v_{j,h,t}^{\tau} - m_t^{\tau}$. Substituting it into the above equation gives:

$$\tilde{v}_{j,h,t}^{\tau} + \varepsilon_{i,j,h,t} > \max_{\{l,k\} \in \mathcal{D}_{\tau,t}} [\tilde{v}_{k,l,t}^{\bar{\tau}} + \varepsilon_{i,k,l,t}] - (m_t^{\tau} - m_t^{\bar{\tau}}) - \text{PMC}_{i,t}^{\bar{\tau}}. \quad (10)$$

The term $(m_t^{\tau} - m_t^{\bar{\tau}})$ captures the decrease in household lifetime utility caused by the reduction in wealth due to financial moving costs, which change the household type from τ to $\bar{\tau}$. Since $(m_t^{\tau} - m_t^{\bar{\tau}})$ is unobserved, we parametrize it as a function of financial moving costs, which depend on household characteristics $\bar{Z}_{i,t}$, homeownership status in the previous and the current period. Formally, we define it as:

$$m_t^{\tau} - m_t^{\bar{\tau}} = \text{FMC}_{i,t}^{\bar{\tau}} \gamma_{i,\text{fmc}}^{\bar{\tau}}. \quad (11)$$

We have four types of financial moving costs: If a renter is buying a house at a price $P_{d_{i,t}}^{\text{buy}}$, where $d_{i,t} = \{j, h\}$ defines the housing product j that is purchased, then $\text{FMC}_{i,t}^{\bar{\tau}} = 0.03 \times P_{d_{i,t}}^{\text{buy}}$.¹⁸ If instead a homeowner is selling its house at price $P_{d_{i,t-1}}^{\text{sell}}$ to buy a new one at

¹⁸Housing transaction costs, both when selling and buying, amount to roughly 3% of the purchase price. In Norway the two dominant forms of ownership are either the “traditional” one or ownership through a co-op or housing association (“borettslag”). In practice, the ownership form has not played a significant role since the 1980s, except for the tax paid to the government after purchasing a property. Traditional properties come with a 2.5% government tax, while co-ops are tax-exempt. However, for two properties of equal characteristics other than the ownership form, we expect the transaction price to be slightly higher for the co-op, as the seller can factor in the tax exemption that the buyer gets. As a result, we set the cost of purchasing a property of either ownership form to be the same at 3%, where 2.5% captures the tax, and additional minor fees add up to the remaining 0.5%. While selling a property is tax-free, broker fees are typically between 1.5% and 3.5% of the transaction price. Given that the extra cost of buying and selling are approximately the same, and we do not have granular data on this extra cost, we set the transaction fee for both at 3%.

price $P_{d_{i,t}}^{\text{buy}}$, then $\text{FMC}_{i,t}^{\bar{\tau}} = 0.03 \times (P_{d_{i,t}}^{\text{buy}} + P_{d_{i,t-1}}^{\text{sell}})$. If a homeowner is selling its house at price $P_{d_{i,t-1}}^{\text{sell}}$ to become a renter, then $\text{FMC}_{i,t}^{\bar{\tau}} = 0.03 \times P_{d_{i,t-1}}^{\text{sell}}$. This captures the idea that selling a property implies facing financial costs similar to purchasing. Last, if, instead, a renter moves to another rented property, then we set the financial moving cost to zero, $\text{FMC}_{i,t}^{\bar{\tau}} = 0$. Hence, $\text{FMC}_{i,t}^{\bar{\tau}}$ has the superscript $\bar{\tau}$, which means that the financial cost depends on the household type. We let $\gamma_{i,\text{fmc}}^{\bar{\tau}} = \bar{Z}'_{i,t} \gamma_{\text{fmc}}$ to allow a marginal change in wealth to have a different impact on household utility depending on household characteristics, and $\text{PMC}_{i,t}^{\bar{\tau}} = \bar{Z}'_{i,t} \gamma_{\text{pmc}}$. The probability that a household stays in its current property of type j with ownership status h in a given period t becomes:

$$\Pr_{\text{stay},i,t}^{\tau,\bar{\tau}} = \frac{\exp(\tilde{v}_{j,h,t}^{\tau})}{\exp(\tilde{v}_{j,h,t}^{\tau}) + \sum_{l=0}^1 \sum_{k=0}^{J_{\bar{\tau}}t} \exp(\tilde{v}_{k,l,t}^{\bar{\tau}} - \text{FMC}_{i,t}^{\bar{\tau}} \gamma_{i,\text{fmc}}^{\bar{\tau}} - \bar{Z}'_{i,t} \gamma_{\text{pmc}})}. \quad (12)$$

The likelihood contribution of each household's sequence of move/stay decisions is denoted $L_i^{\text{stay}}(\tilde{v}, \gamma_{\text{fmc}}, \gamma_{\text{pmc}})$ and is given by:

$$L_i^{\text{stay}}(\tilde{v}, \gamma_{\text{fmc}}, \gamma_{\text{pmc}}) = \prod_{t=t_{0,i}}^{t_{0,i}+T_i} (P_{\text{stay},i,t}^{\tau,\bar{\tau}})^{\mathbb{1}_{[d_{i,t}=d_{i,t-1}]}} (1 - P_{\text{stay},i,t}^{\tau,\bar{\tau}})^{\mathbb{1}_{[d_{i,t} \neq d_{i,t-1}]}}. \quad (13)$$

3.3.2 Real Estate Investors

We now characterize real estate investors' decisions. While regular households can stay or move to another property, either purchased or rented, investors do not move but just buy and sell. Moreover, investors are financially unconstrained, so they never change their type.

We assume the demand for product j at time t is an exponential function of house prices $P_{j,t}$ and latent demand $\delta_{j,t}^s$:

$$\mathcal{S}_{j,t}^s = \exp\{\alpha P_{j,t} + \delta_{j,t}^s\}, \quad (14)$$

where α is the price elasticity of demand, and the remaining part, $\delta_{j,t}^s$, reflects latent demand for house product j unrelated to the price, and $\mathcal{S}_{j,t}^s$ is the stock of housing for product j held by investors at time t . Taking the logarithm of equation (14), the demand function becomes:

$$\ln(\mathcal{S}_{j,t}^s) = \alpha P_{j,t} + \delta_{j,t}^s. \quad (15)$$

3.4 Market Clearing

We now define the market clearing condition we use in the counterfactuals to determine the new equilibrium property prices. Our model delivers predictions on households' housing demand and supply and the housing portfolio held by investors.

For each product j at time t , the market clears when the housing supply of that product equals demand, that is:

$$\underbrace{\# \text{New Houses} + \# \text{Houses net supplied by investors} + \# \text{Houses sold by owners}}_{\text{Supply}} = \underbrace{\# \text{Houses bought by renters} + \# \text{Houses bought by switchers} + \# \text{Houses bought by entries}}_{\text{Demand}}.$$

In the demand part, switchers are homeowners who change residency, and entries are new households that enter the sample. Based on households' and investors' decision probabilities, we can construct a market clearing condition that determines the equilibrium price $P_{j,t}$ of each housing product j at time t as:

$$\begin{aligned} & \underbrace{\mathcal{S}_{j,t}^{New} + (\mathcal{S}_{j,t-1}^s - \mathcal{S}_{j,t}^s) + \sum_{\tau} \sum_{i \in \tau} (1 - \Pr_{i,t}^{\tau, \bar{\tau}}(\{j, 1\} | \{j, 1\}))}_{\text{Supply}} \quad (16) \\ = & \underbrace{\sum_{\tau} \sum_k \sum_{i \in \tau} \Pr_{i,t}^{\tau, \bar{\tau}}(\{j, 1\} | \{k, 0\}) + \sum_{\tau} \sum_k \sum_{i \in \tau} \Pr_{i,t}^{\tau, \bar{\tau}}(\{j, 1\} | \{k, 1\}) + \sum_{\bar{\tau}} \mathcal{N}_t^{\bar{\tau}} \Pr_{j,h=1,t}^{\bar{\tau}}}_{\text{Demand}}, \end{aligned}$$

Here $\mathcal{S}_{j,t}^{New}$ is the number of newly constructed houses of type j in year t , and $\mathcal{S}_{j,t-1}^s - \mathcal{S}_{j,t}^s$ is the net supply of type j properties by investors at time t . The probability $\Pr_{i,t}^{\tau, \bar{\tau}}(d_{i,t} | d_{i,t-1})$ indicates the probability that household i with previous decision $d_{i,t-1}$ makes decision $d_{i,t}$ at time t , changing its household type from τ to $\bar{\tau}$ with this decision. The supply is determined by the number of new houses, the net supply of investors, and the number of properties sold by homeowners. Together, this quantity corresponds to the last term in the supply (left-hand side) part of equation (16). More precisely, we sum up the probabilities that each homeowner belonging to type τ and owning product j does not remain in its property $1 - \Pr_{i,t}^{\tau, \bar{\tau}}(\{j, 1\} | \{j, 1\})$; which gives the number of type τ homeowners that sell product j at time t . Next, to obtain the aggregate supply of owned housing product j at time t , we sum across all types τ .

Demand for each product j at time t is determined by the number of housing products j purchased by renters plus that purchased by homeowners and households moving to Oslo, which corresponds to the first, the second, and the third term in the demand (right-hand side) part of equation (16). For the first two terms, summing up the probabilities that each household belonging to type τ and renting (or owning) product k chooses to purchase

product j gives the number of products j bought by renters (or homeowners) of type τ at time t who used to live in product k . To obtain the total housing demand, we sum over all property types that a household used to live in and over all household types. For the last term, $\Pr_{j,h=1,t}^{\bar{\tau}}$ is the probability that a household chooses to own type j conditional on the decision of moving, defined in equation (5); $\mathcal{N}_t^{\bar{\tau}}$ is the number of households of type $\bar{\tau}$ that enter into Oslo at time t . Summing $\mathcal{N}_t^{\bar{\tau}} \Pr_{j,h=1,t}^{\bar{\tau}}$ over all household types gives us the number of product j bought by new entries.

Formally, the decision probabilities are obtained as:

$$\Pr_{i,t}^{\tau,\bar{\tau}}(d_{i,t} = \{j, h\} \mid d_{i,t-1}) = \frac{\exp(\tilde{v}_{j,h,t}^{\bar{\tau}} - \text{FMC}_{i,t}^{\bar{\tau}} \hat{\gamma}_{i,\text{fmc}}^{\bar{\tau}} - \bar{Z}'_{i,t} \hat{\gamma}_{\text{pmc}})}{\exp(\tilde{v}_{0,t}^{\bar{\tau}}) + \sum_{l=0}^1 \sum_{k=0}^{J_{\bar{\tau}t}} \exp(\tilde{v}_{k,l,t}^{\bar{\tau}} - \text{FMC}_{i,t}^{\bar{\tau}} \hat{\gamma}_{i,\text{fmc}}^{\bar{\tau}} - \bar{Z}'_{i,t} \hat{\gamma}_{\text{pmc}})}. \quad (17)$$

Note that the financial moving costs $\text{FMC}_{i,t}^{\bar{\tau}}$ is determined by $d_{i,t}$ and $d_{i,t-1}$, as it is a function of the property price of the product that will be purchased or sold. The price that clears each j, t combination will determine the value of the housing asset for sellers and buyers. We do not require the rental market to clear, as renters can always consume less housing by living together. Investors own all the houses available for rent. Renters are price-takers, and rental prices are based on a fixed rent-to-price ratio.

4 Estimation

We estimate the model in four steps. First, we estimate \tilde{v} according to households' housing product choices. Second, we estimate $\gamma_{\text{fmc}}, \gamma_{\text{pmc}}$ from the decisions to stay or move taking \tilde{v} as given. Third, we recover the determinants of households' flow utility. Last, we estimate the parameters of the investors' portfolio problem.

4.1 First Stage: Value Functions

The closed-form solution for maximizing the likelihood of choosing a housing product (including the outside option), conditional on moving, is $\sum_{i=1}^N \left(\ln(L_i^{\text{product}}(\tilde{v})) + \ln(L_i^{\text{out}}(\tilde{v})) \right)$, is given by:

$$\hat{\tilde{v}}_{j,h,t}^{\bar{\tau}} = \ln(\widehat{\Pr}_{j,h,t}^{\bar{\tau}}) - \frac{1}{2J_{\bar{\tau}t} + 1} \left(\ln(\widehat{\Pr}_{0,t}^{\bar{\tau}}) + \sum_{l=0}^1 \sum_{k=1}^{J_{\bar{\tau}}} \ln(\widehat{\Pr}_{k,l,t}^{\bar{\tau}}) \right), \quad (18)$$

where $\widehat{\Pr}_{j,h,t}^{\bar{\tau}}$ is the empirical probability that a household of type $\bar{\tau}$ chose housing product j with ownership status h at time t conditional on moving. Instead of using these observed probabilities directly, we use the kernel smoothing method of Bayer et al. (2016) to account for the product choice decisions of similar household types. The kernel assigns weights across household types depending on how similar the households are, measured by their characteristics. The benefit of the kernel smoothing method is that it mitigates potential small sample problems caused by having many household types.¹⁹

4.2 Second Stage: Moving Costs

Once we recover \widehat{v} , we find $\gamma_{\text{fmc}}, \gamma_{\text{pmc}}$ that maximize the likelihood of the decision to stay and move, taking \widehat{v} as given, that is :

$$\max_{\gamma_{\text{fmc}}, \gamma_{\text{pmc}}} \sum_i^N \ln(L_i^{\text{stay}}(\widehat{v}, \gamma_{\text{fmc}}, \gamma_{\text{pmc}})). \quad (21)$$

The estimated $\widehat{\gamma}_{\text{fmc}}, \widehat{\gamma}_{\text{pmc}}$ can be used to recover the true choice-specific value functions $v_{j,h,t}^{\tau} = \widetilde{v}_{j,h,t}^{\tau} + m_t^{\tau}$. Notice that $m_t^{\tau} - m_t^{\bar{\tau}}$ captures how the decrease in wealth due to financial moving costs affects the lifetime utility of household type τ . If we now normalize the average utility of households without wealth to zero, then the impact of wealth on households' utility can be recovered by multiplying household wealth with the marginal utility of wealth. Thus, we set $m_t^{\tau} = A_{i,t}^{\tau} \gamma_{i,\text{fmc}}^{\tau}$, where i is the median income household of type τ . This recovers the wealth effect on household utility $v_{j,h,t}^{\tau}$ for all household types.

¹⁹Specifically, we calculate the empirical probability of a type $\bar{\tau}$ household choosing housing product j with ownership h in time t conditional on moving, taking into account all household types that made the same housing decision:

$$\widehat{\Pr}_{j,h,t}^{\bar{\tau}} = \frac{\sum_{i=1}^N \mathbb{1}_{[d_{i,t}=\{j,h\}]} \cdot W^{\bar{\tau}}(\bar{Z}_{i,t})}{\sum_{i=1}^N W^{\bar{\tau}}(\bar{Z}_{i,t})}, \quad (19)$$

where $W^{\bar{\tau}}(\bar{Z}_{i,t})$ is the weight assigned to household i with characteristics $\bar{Z}_{i,t}$. We assign higher weights to household types with higher similarity to $\bar{\tau}$ in the household characteristic space. The weight is the product of L normal kernels N :

$$W^{\bar{\tau}}(\bar{Z}_{i,t}) = \prod_{l=1}^L \frac{1}{b_l^{\bar{\tau}}} N\left(\frac{\bar{Z}_{i,t}(l) - \bar{Z}^{\bar{\tau}}(l)}{b_l^{\bar{\tau}}}\right), \quad (20)$$

where L is the dimension of Z , $Z(l)$ is l th attribute of household characteristics, and $b_l^{\bar{\tau}}$ is the bandwidth of the l th attribute determined by cross validation.

4.3 Third Stage: Indirect Utility

To recover the per-period utility, we need to define the transition dynamics of the state variables. We model the transition of the choice-specific value functions $v_{j,h,t}^\tau$ and house prices $P_{j,t}$ as:

$$v_{j,h,t}^\tau = \psi_{0,j,h}^\tau + \sum_{l=1}^2 \psi_{1,l} v_{j,h,t-l}^\tau + \sum_{l=1}^2 \psi_{2,l} P_{j,t-l} + \psi_{3,j,h}^\tau t + \omega_{j,h,t}^\tau, \quad (22)$$

$$P_{j,t} = \phi_{0,j} + \sum_{l=1}^2 \phi_{1,l} P_{j,t-l} + \phi_{2,j} t + \rho_{j,t}. \quad (23)$$

Knowing $v_{j,h,t}^\tau$, γ_{pmc} , γ_{fmc} , and the transition probabilities allows us to calculate the mean flow utility for each type and product, $u_{j,h,t}^\tau$, according to:

$$u_{j,h,t}^\tau = v_{j,h,t}^\tau - \beta \mathbb{E} \left[\ln \left(e^{v_{j,h,t+1}^\tau} + \sum_{l=0}^1 \sum_{k=0}^{J_\tau} e^{v_{k,l,t+1}^\tau - \text{FMC}_{i,t+1}^\tau \hat{\gamma}_{i,\text{fmc}}^\tau - \bar{Z}'_{i,t+1} \hat{\gamma}_{\text{pmc}}^\tau} \right) \mid s_{i,t}, d_{i,t} = \{j, h\} \right], \quad (24)$$

where, in practice, $s_{i,t}$ includes all the variables on the right-hand side of equations (22) and (23), and β is set to 0.95. For each type, τ , product, j , ownership status h , and time, t , we now have the necessary information to simulate the expectation on the right-hand side of equation (24). To do this, we draw a large number of $v_{j,t+1}$ and $P_{j,t+1}$ from their empirical distributions. Specifically, we index each random draw with r , and produce these variables by sampling from the empirical distribution of errors, $\omega_{j,h,t}^\tau$ and $\rho_{j,t}$, obtained when estimating each of these processes respectively. For the other variables we use the observed values of the current states. The draws of house prices are used to determine housing wealth, which determines households' type in the next period τ_{t+1} . For each draw r , we calculate a per-period flow utility $u_{j,t}^\tau$ using equation (24). The average across draws is the simulated $u_{j,t}^\tau$.

To estimate the determinants of flow utility of household type τ choosing housing product j with ownership status h at time t , we follow Bayer et al. (2016) and define:

$$u_{j,h,t}^\tau = \alpha_0^\tau + \alpha_h^\tau + \alpha_t^\tau + X_{j,t} \alpha_x^\tau + \alpha_r^\tau R_{j,t} + \xi_{j,h,t}^\tau. \quad (25)$$

The unobserved house and neighborhood attributes ($\xi_{j,h,t}^\tau$) impact the flow utility of different household types differently. The parameters α_0^τ , α_h^τ , α_t^τ represent a household-type-specific constant, the utility of homeownership, and year fixed effects. The parameter α_x^τ captures the impact of observed housing attributes on household utility. The parameter α_r^τ

captures the disutility of user cost of housing $R_{j,t}$, which includes the rental payments.²⁰ To address the endogeneity of prices, we exploit the estimated $\widehat{\gamma}_{\text{fmc}}^\tau$ that reflects the marginal utility of wealth. Assuming that households have the same marginal utility of wealth as that of income, we can rewrite equation (25) as:

$$u_{j,h,t}^\tau + \widehat{\gamma}_{\text{fmc}}^\tau R_{j,t} = \alpha_0^\tau + \alpha_h^\tau + \alpha_t^\tau + X_{j,t} \alpha_x^\tau + \xi_{j,h,t}^\tau. \quad (26)$$

We estimate the parameters $\alpha_0^\tau, \alpha_h^\tau, \alpha_t^\tau, \alpha_x^\tau$ with a linear regression.

4.4 Fourth Stage: Investors' Portfolio Choice

We estimate α and $\delta_{j,t}^s$ in the investors' model such that: (i) the investors' demand function defined by equation (15) holds, (ii) households' decisions $\Pr_{i,t}^{\tau, \bar{\tau}}(d_{i,t} | d_{i,t-1})$ defined by equation (17) hold, and (iii) the market clearing conditions defined by equation (16) are satisfied. To achieve this, we first normalize the latent demand for the first house product to zero, which ensures the identification of the remaining parameters.

We then guess an α and a vector of $\delta_{j,t}$ that satisfy conditions (i) to (iii) with a new set of market clearing prices. We compare the new market clearing prices with observed prices and then update α and $\delta_{j,t}$ until the model-implied market clearing prices closely match the observed house prices. To evaluate the robustness of the iterative procedure, we verify that the final parameter estimates are independent of the starting values. The resultant investor demand function is the one we use in all counterfactuals.

5 Results

5.1 Model Estimates

To estimate the financial and psychological moving cost parameters γ_{fmc} and γ_{pmc} , we use the likelihood function (13), which is based on households' mobility choices. Table 4 reports the results. For the financial moving cost, the constant represents the average marginal disutility from 1,000 NOK (~ 80 USD) of financial moving costs. We allow financial moving costs to depend on income. The income variable is the median value for each of the five groups. Psychological moving cost depends on households' income, age group, and family size.²¹ The takeaways from Table 4 are as follows. First, and as expected, financial moving costs

²⁰We assume that even owners pay a user cost of housing that can be proxied by the rental price of an equivalent property. User cost for owners aim to capture mortgage payments, property taxes, and any maintenance and renovation costs.

²¹The characteristics are defined in footnote 13.

matter less for high-income households. Second, psychological moving costs also decrease with household income. Third, older households face high psychological moving costs, and couples higher than singles.

Table 4 FINANCIAL AND PSYCHOLOGICAL MOVING COST

	Estimate
<i>Financial Moving Cost</i>	
Constant	0.01*** (0.000)
Income	-0.01*** (0.000)
<i>Psychological Moving Cost</i>	
Constant	6.56*** (0.01)
Income	-0.24*** (0.02)
Age: 35 - 54	1.03*** (0.01)
Age: 55 +	1.85*** (0.01)
Couple	0.31*** (0.01)

Notes: This table presents the estimated parameters of financial moving costs γ_{fmc} and psychological moving costs γ_{pmc} . *Income* is the median income (in millions of NOK) of each income group. *Age: 35-54* is a dummy variable taking the value of one if the head of a household is between 35 and 54 years old, and zero otherwise. *Age: 55+* is a dummy variable taking the value of one if the head of a household is above 55 years old, and zero otherwise. *Couple* is a dummy variable if a household is a couple, and zero otherwise. Bootstrapped standard errors are presented in brackets.

5.2 Decomposition of Flow Utilities

We now regress the flow utility $u_{j,h,t}^\tau$ on property and neighborhood characteristics, excluding user costs in line with equation (26). The coefficients presented in Table 5 reveal the drivers of households' flow utility. Column (1) shows that utility is increasing in neighborhood quality, and homeownership is preferred to renting. In column (2), we interact neighborhood quality and property size with homeownership. Homeowners have weaker preferences for high quality relative to renters, presumably because renting in high-quality districts is more affordable in the short run relative to owning, but stronger preferences for larger property sizes. In column (3), we interact neighborhood quality, property size, and homeownership with household income. This is the specification we use for our willingness to pay analysis, as it allows us to illustrate how the willingness to pay for housing attributes varies across

different income groups and helps us capture the distributional effects in our counterfactuals. We find that households’ utility for high-quality neighborhoods is increasing in income. Low-income households prefer larger properties, and homeownership delivers higher utility than renting at an increasing rate with income.

Panel B of Table 6 reports households’ willingness to pay for housing attributes by income. We scale the estimates by dividing the estimated coefficients from equation (26) by the marginal utility of wealth $\hat{\gamma}_{\text{fmc}}^{\tau}$. These ratios allow us to interpret the magnitudes of the determinants of the flow utility. In addition, we report some household characteristics in Panel A, to guide the interpretation of the willingness to pay, and the ratio of willingness to pay to median income in Panel C.

This willingness to pay for neighborhood quality is increasing in income. Except for the lowest income group, all households willingly pay more for living in high-quality neighborhoods than middle-quality ones. For example, an average household with an annual income between 600,000 and 800,000 NOK ($\sim 48,000$ - $64,000$ USD) is willing to pay annually 14,600 NOK ($\sim 1,168$ USD) to live in a high-quality neighborhood relative to a middle-quality.

Regarding house size and homeownership, we find that households in the three lowest income groups are willing to pay increasingly more for larger properties. Those in the fourth income group pay the highest for middle-sized properties. Top-income groups’ willingness to pay is decreasing in property size. We interpret this as a reflection of family composition, as families with young children, who need more living space, often earn less than older households whose children have left the house and now need less space. The willingness to pay for homeownership is increasing across income groups.²²

6 Counterfactuals

We simulate three counterfactual scenarios to quantify the distributional effects of leverage limits. First, we simulate a scenario for 2017 in which the LTI is three instead of five. We focus on the counterfactual house prices and mobility patterns across the income distribution. Second, to understand how our counterfactuals interact with households’ preferences, we compare the effect of the change in LTI to a scenario where preferences of households in the first four income groups are shifted closer to those of the top income group, reducing the distance between them by 10%. Simulating this change is a way to quantify the importance of preferences in residential choices and proxy for housing policies, such as schooling vouchers

²²These results are comparable to the willingness-to-pay estimates of Bayer et al. (2016), who show that a household with average income is willing to pay up to 2,256 USD for a 10% increase in amenities in their neighborhood.

Table 5 DETERMINANTS OF FLOW UTILITY

	Flow Utility \hat{u}		
	(1)	(2)	(3)
High Quality	0.07*** (0.03)	0.09*** (0.03)	-0.08** (0.04)
Low Quality	-0.11*** (0.04)	-0.11*** (0.04)	0.04 (0.03)
Three Rooms	0.06 (0.04)	0.04 (0.05)	0.29*** (0.05)
Four Rooms and Above	-0.05 (0.04)	-0.10** (0.05)	0.43*** (0.07)
Home-Ownership	0.11*** (0.01)	0.07*** (0.02)	-0.05*** (0.02)
High Quality \times Home-Ownership		-0.05*** (0.02)	
Low Quality \times Home-Ownership		0.01 (0.02)	
Three Rooms \times Home-Ownership		0.05** (0.02)	
Four Rooms and Above \times Home-Ownership		0.11*** (0.03)	
High Quality \times Income			0.21*** (0.03)
Low Quality \times Income			-0.21*** (0.05)
Three Rooms \times Income			-0.33*** (0.03)
Four Rooms and Above \times Income			-0.67*** (0.08)
Home-Ownership \times Income			0.22*** (0.03)
Year FE	Yes	Yes	Yes
Household Type FE	Yes	Yes	Yes
Observations	81,510	81,510	81,510
Adjusted R ²	0.83	0.83	0.84

Notes: This table reports the determinants of estimated flow utility. *High Quality* and *Low Quality* are dummy variables, with the omitted category being *Middle Quality*. *Three Rooms* is a dummy variable equal to one if a property type has three bedrooms, and zero otherwise. *Four Rooms and Above* is a dummy variable equal to one if a property type has four or more bedrooms, and zero otherwise. *Home-Ownership* is a dummy variable that equals one for homeowners and zero for renters. *Income* is each income group's median Income (in millions of NOK).

Table 6 WILLINGNESS TO PAY FOR HOUSING ATTRIBUTES

	Income Groups				
	1 <400k	2 400-600k	3 600-800k	4 800-1,000k	5 >1,000k
<i>A. Household Characteristics</i>					
Home-Owners	28.0%	41.5%	53.6%	62.2%	70.0%
Young	36.5%	34.9%	26.1%	17.6%	10.1%
<i>B. Willingness to Pay</i>					
High Quality	-3.3	4.0	14.6	34.6	170.1
Low Quality	-2.8	-11.5	-24.1	-47.9	-209.0
Three Rooms	30.6	25.0	16.7	1.3	-103.6
Four Rooms and Above	36.2	19.0	-6.2	-53.6	-374.9
Ownership	2.25	10.87	23.47	47.18	207.87
<i>C. Willingness to Pay to Income</i>					
High Quality	-1.1%	0.8%	2.2%	3.9%	13.9%
Low Quality	-1.0%	-2.3%	-3.6%	-5.4%	-17.1%
Three Rooms	10.6%	5.1%	2.5%	0.1%	-8.5%
Four Rooms and Above	12.6%	3.9%	-0.9%	-6.1%	-30.7%
Ownership	0.8%	2.2%	3.5%	5.4%	17.0%

Notes: This table reports household characteristics and the average willingness to pay for housing attributes of households with different income levels. Panel A reports the percentage of households that are homeowners (Home-Owners), below 35 years old (Young), and couples (Couple). In Panel B, the first row suggests how much more an average household is willing to pay (in thousands of NOK) for being in a High-Quality district relative to a Middle-Quality district every year. The second row suggests how much more an average household is willing to pay (in thousands of NOK) for living in a Low-Quality district relative to a Middle-Quality district every year. The third (fourth) row suggests how much an average household in the income group is willing to pay (in thousands of NOK) for living in an apartment with three rooms (more than four rooms) compared with living in a one- or two-room apartment. The last row shows the average willingness to pay for living in their own house compared to a rented house. Panel C reports the willingness to pay in Panel B as a fraction of the median income of each income group.

and public housing assistance programs.²³ Third, we simulate a policy that curbs investors' incentive to transact in housing markets, for example, due to increased taxation on second homes or rental income. In the model, this corresponds to making investors more price inelastic. As investors mostly own properties in low-quality districts, this counterfactual sheds light on whether policies that make real estate investors more inelastic can incentivize low-income households to move into high-quality neighborhoods.

Before showing the counterfactual results, we present in Figure 4 the 2017 baseline distribution, by homeownership status, of households across house quality and income. The

²³These programs encourage low-income households to become homeowners and live in high-quality neighborhoods.

baseline distribution provides a benchmark for the counterfactuals. The left panel presents the “stocks”, that is the distribution of all households. The right panel shows the “flows”, which corresponds to the distribution of moving households. The “flow” estimates represent the model’s predicted probability of moving into a specific property type times the likelihood of moving. For both figures, the sum of the vertical bars in a row between owners and renters adds up to 100%. Note that the total number of households across income groups is unequal. The two lowest income groups represent roughly 50% of households, and the top income group contains the fewest households. This uneven distribution mimics the skewed income distribution, to capture the increasingly different preferences that households in the top income groups have relative to the bottom ones.

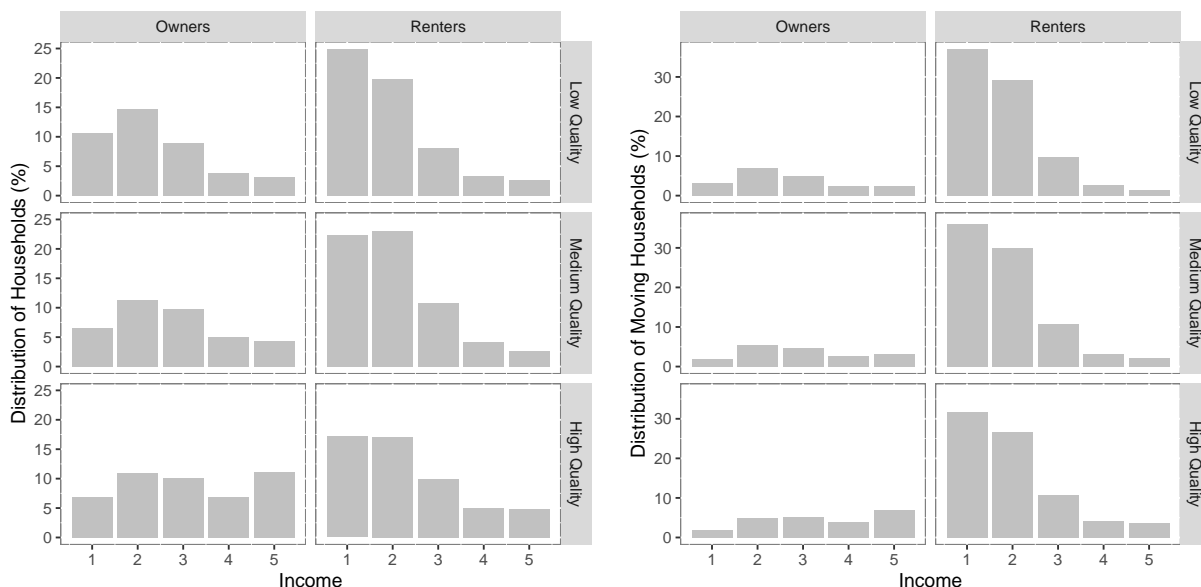
There are three takeaways from the left figure. First, across all quality levels, households in the first two income groups are always more likely to be renters than owners, whereas the opposite is true for the two top income groups. This suggests low-income households are primarily young people who, for instance, have not yet accumulated enough wealth and income to afford a house. Second, the proportion of owners relative to renters increases with district quality. This means the real estate investors’ portfolios are tilted towards properties in low-quality districts. Third, as expected, the largest proportion of high-income households is in high-quality districts, and low-income households dominate low-quality districts.

Similar patterns arise from the right figure, which focuses on movers. In addition, it presents two new insights. First, movers are more likely to end up as renters, partly due to renters’ lower cost of moving. Second, low-income households are much more mobile than high-income ones. The high mobility of low-income households mainly reflects that those are unsettled young households.

6.1 More Stringent LTI Limit

Table 7 shows the results of our counterfactual LTI change relative to the baseline. For every row of the table, we present shares or probabilities across five income groups, both for the baseline level of LTI (labeled as *Base*) and for the percentage change between counterfactual and baseline (labeled as Δ). The top panel shows outcomes for all households in our data, while the bottom panel shows outcomes for households who moved in 2017.

Starting from the top panel, our model shows that for the baseline level of LTI, homeowners in the lowest (highest) income group have access to 52.9% (73.1%) of all 45 housing products. As the LTI becomes more stringent, low (high) income households’ choice set shrinks, as they lose access to 14.3% (1%) of housing products. The lowest and highest income groups are only marginally affected by a lower LTI in the share of properties they



Notes: These figures plot the distribution of all households (left panel) and of moving households (right panel) across housing quality levels (low quality at the top, middle quality in the middle, high quality at the bottom) by households' income groups and across owners and renters.

Figure 4 STOCKS AND FLOWS OF HOUSEHOLDS ACROSS QUALITY AND INCOME

can access but for different reasons. High-income households are unaffected because the LTI constraints are mostly not binding. In contrast, low-income households are affected to a small extent because several houses were already unaffordable even with a higher LTI. The largest effect of the LTI limit on the housing choice set is for households in the middle of the income distribution. For example, the third income group had access to 68.4% of the housing market before the introduction of the LTI, which then dropped by 17.5%.

Moreover, in the baseline 14.2% (22.9%) of housing products are of high quality for households in the lowest (highest) income group. In the counterfactual, their share of high-quality products drops by 12.5% (3.2%). Again, the lowest and highest income groups are least affected by a lower LTI limit for the same reason as earlier. Middle-income households experience a drop in their access to high-quality housing of 26.2% as a result of a tighter LTI.

The other two rows of the top panel of Table 7 quantify homeowners and renters' mobility patterns across the income distribution. Three results merit attention. First, as expected, renters are more likely to move than homeowners. Second, high-income homeowners are more likely to move than low-income ones. Third, a stricter LTI limit almost does not affect homeowners' probability of moving. The reason for this is just that the baseline probability of moving is very low.

The second panel of Table 7 focuses on households that move. The probability of renters

becoming homeowners, mostly representing first-time buyers, is 5.5% (44.1%) for the lowest (highest) income group. A tighter LTI limit reduces it by 11.6% (0%). Among movers, the probability of climbing the housing quality ladder is unaffected by LTI change for the lowest and highest income groups. In contrast, lower LTI reduces the likelihood of moving from a low-quality to a high-quality district by 1.3% for middle-income households.²⁴

Table 7 EFFECT OF CHANGE IN LTI LIMIT ON CHOICE SETS AND MOVING PROBABILITIES

	LTI	Income Groups				
		1	2	3	4	5
All households						
Owners' Share of Total Products in Choice Set	Base	52.9%	61.3%	68.4%	73.3%	73.1%
	Δ	-14.3%	-18.8%	-17.5%	-15.2%	-1.0%
Owners' Share of High Quality Products in Choice Set	Base	14.2%	16.0%	18.7%	21.3%	22.9%
	Δ	-12.5%	-19.4%	-26.2%	-20.8%	-3.2%
Owners' Moving Probability	Base	4.3%	4.0%	3.9%	4.2%	5.9%
	Δ	0.0%	0.0%	-0.3%	-0.4%	0.0%
Renters' Moving Probability	Base	16.5%	13.8%	11.5%	11.1%	13.5%
	Δ	-0.4%	-1.2%	-0.9%	-0.9%	0.0%
Movers						
From Renting to Owning	Base	5.5%	15.2%	27.7%	39.0%	44.1%
	Δ	-11.6%	-10.6%	-3.6%	-1.9%	0.0%
From Low to High Quality	Base	10.4%	13.2%	15.1%	17.7%	16.5%
	Δ	-0.7%	-0.1%	-1.2%	-1.3%	0.0%

Notes: This table reports the mobility of households for each income quintile under the baseline scenario (Base) and the percentage changes under the counterfactual scenario of changing LTI limits (Δ). The first panel reports the share of available products for home purchasing (Owners' Share of Total Products in Choice Set), the share of available high-quality products for home purchasing (Owners' Share of High-Quality Products in Choice Sets), the probability of moving for homeowners (Owners' Moving Probability), and the probability of moving for renters (Renters' Moving Probability). The second panel considers only movers, where the probability of movers changing from renting to owning (From Renting to Owning) and the probability of movers changing from low-quality districts to high-quality districts (From Low to High Quality) are reported.

²⁴In Appendix C we report the same table as Table 7 for a longer-term scenario. More specifically, we look at the effects of the same change in the LTI limit 2017, allowing the change to occur since 2016. The results are very close to our 1-year baseline, implying that the effects of LTI changes are likely permanent and exhibit no decay after one year.

6.2 Restrictions on Investors

In Table 8, we report changes in equilibrium prices between the baseline and the counterfactual LTI under two scenarios. In the first scenario, we use investors’ price elasticity of demand in our baseline. In the second scenario, we reduce investors’ price elasticity of demand by 50% relative to the baseline. As mentioned earlier, such change can be microfounded by, for example, an increase in transaction fees, which increases the cost of portfolio rebalancing.

A more stringent LTI reduces prices by up to 2.6%. Large and high-quality properties drop the most in value. This happens because those are the most expensive properties. As affordability constraints tighten, those houses drop from middle-income households’ choice sets, as displayed in the top panel of Table 7. As middle-income households shift their demand towards less expensive houses—smaller properties and homes in lower-quality districts—the price impact of stricter LTI limits on this housing segment becomes small.

We now make investors more price inelastic, which has a heterogeneous price impact across size and quality distribution. In particular, the prices of small and high-quality houses increase by 9.1% and 9.7%. In stark contrast, large property prices drop by 0.8%. We interpret these results as follows. Inelastic investors are less responsive to price changes. As a result, households wishing to buy properties from investors must pay relatively more. High-income households are the least price-sensitive and have the greatest willingness to pay for high-quality homes and smaller properties, and therefore, will pay the “premium” to move. More price-sensitive households will instead mainly respond to the new situation of a less liquid housing market by changing their housing demand. Hence, policies increasing the cost of investors to trade will have a heterogeneous impact on house prices (and household wealth) and mobility. The last column of Table 8 shows how prices are affected by a tighter LTI limit when investors are more inelastic. Results are qualitatively in line with the case of elastic investors, but price changes are slightly larger.

6.3 Role of Preferences

The last counterfactual we run aims to quantify the importance of preferences in residential choices. To do so, we adjust all households’ preferences to be closer to those of the top income group. We simulate a scenario where we reduce the distance between preferences of the bottom four income groups and the top one by 10%.²⁵ The experiment proxies for housing policies that encourage low-income households to become homeowners or move to high-quality neighborhoods. In Figure 5 we compare the effects of this counterfactual, labeled

²⁵We obtain the preference disparity between the bottom four income groups and the top income group by calculating the difference in flow utility of the top income group compared to the other four groups, while controlling for identical wealth, family size, and age group.

Table 8 EFFECT OF CHANGE IN LTI LIMIT AND INVESTORS' ELASTICITY ON PRICES

	Elastic investors (base)		Inelastic investors	
	LTI = 5	Δ LTI	LTI = 5	Δ LTI
House Size				
Small	3.2	-0.4%	3.5 (+9.1%)	-0.7%
Medium	4.2	0.0%	4.5 (+6.8%)	-0.3%
Large	6.6	-2.6%	6.6 (-0.8%)	-2.8%
District Quality				
Low	3.0	-0.5%	3.2 (+6.6%)	-0.9%
Medium	3.4	-0.4%	3.6 (+6.6%)	-0.6%
High	4.1	-1.3%	4.5 (+9.7%)	-1.6%

Notes: This table reports the average house prices (in millions NOK) in 2017 under investors' baseline price elasticity (*Elastic investors*) and for a counterfactual in which the price elasticity of investors is lowered by 50% relative to baseline. (*Inelastic investors*). The second and fourth columns report prices under the baseline LTI of 5, while the third and fifth columns report the percentage change in prices when the LTI is reduced to 3. The fourth column also reports the percentage change in prices between scenarios with inelastic and elastic investors. The first panel reports the average house prices for small (1-2 bedrooms), medium (3 bedrooms), and large (4 bedrooms and above), while the second panel is for low, medium, and high-quality districts.

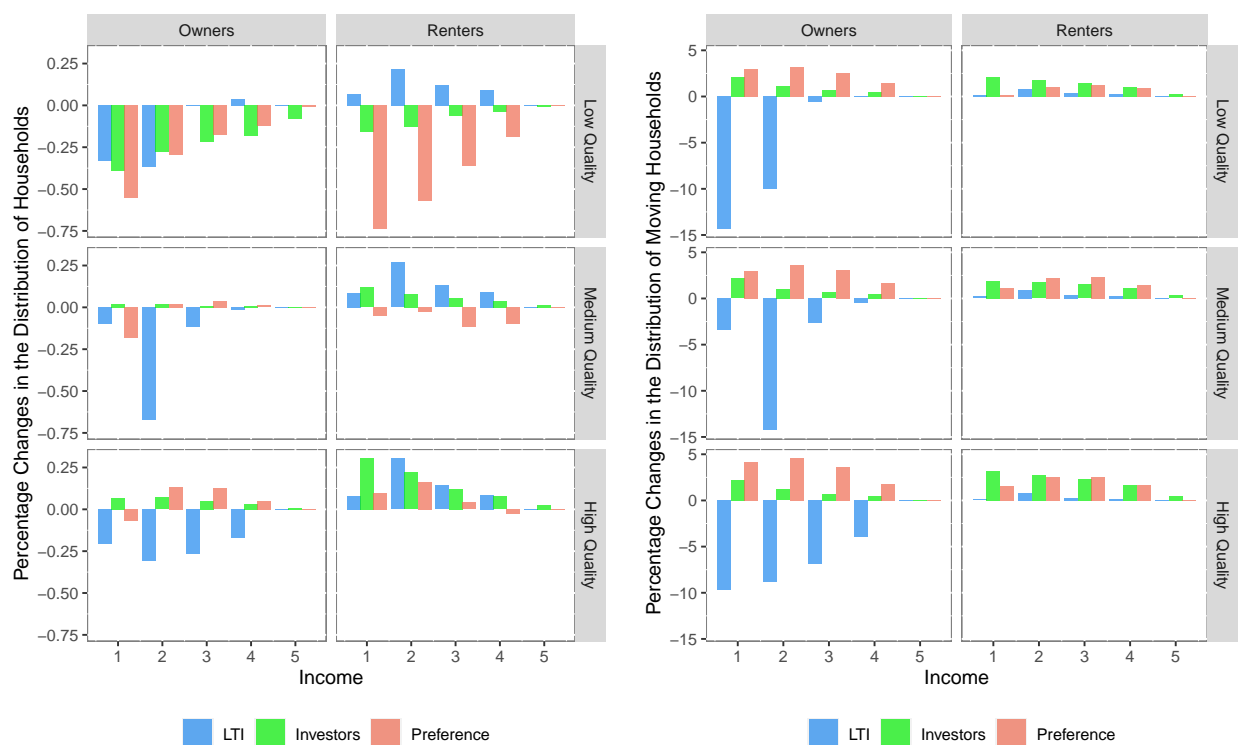
as "Preference," to the effects of the other two counterfactuals presented above, labeled as "LTI" and "Investors". We show counterfactual results for the distributions of households' stocks and flows of properties and present the percentage changes relative to the baseline in Figure 4.

In line with the results in Table 7, both figures show how tighter LTI limits reduce the extent of homeownership and increase rentals across all quality levels, both for movers and for all households. For the case of movers, there is a stronger reduction in ownership relative to the increase in rentals, driven by the larger fraction of households who decide not to move due to the tighter limits. The effect of a lower LTI limit is stronger for low income households, highlighting its regressive effect.

The counterfactual of inelastic investors generates two main effects. First, as the right panel of Figure 5 shows, there is an increase in mobility across quality levels for both owners and renters. This happens because inelastic investors trade less in housing markets, resulting in households substituting for some of those missing real estate trades. Second, as the left panel of Figure 5 shows, there is a reduction in owners and renters in low-quality districts and an increase of mostly renters in medium- and high-quality neighborhoods. This is because investors mostly hold properties in low-quality districts, and becoming inelastic makes them less likely to supply these houses. Consequently, households that were going to reside in low-quality districts are now more likely to rent in a high-quality one.

We find two main effects when we shift households' preferences closer to those of the top income group. First, as the right panel of Figure 5 shows, the likelihood of moving increases across the income distribution, with more substantial changes for low-income households and across quality levels. As expected, the effects are the largest for high-quality districts. Second, as the left panel of Figure 5 shows, households are reallocated from low- to high-quality neighborhoods, both for owners and renters. These effects materialize because the change in preferences increases the benefit of being homeowners and living in high-quality neighborhoods.

The overall takeaway of this comparison is the following. On the one hand, more stringent LTI limits can have regressive effects and increase segregation, as they reduce homeowner-ship rates primarily for low-income households. On the other hand, influencing households' preferences or limiting investor trading in housing markets can have a countervailing effect, as both types of interventions encourage household mobility, especially for low-income households, and create an incentive to move to high-quality neighborhoods.



Notes: These figures plot the percentage changes in the distribution of all households (left panel) and of moving households (right panel) across housing quality levels (low quality at the top, middle quality in the middle, high quality at the bottom) by households' income groups, and across owners and renters. These changes are presented for three different counterfactuals. The blue bars represent the case of tighter LTI limits, the green bars represent the case of inelastic investors, and the orange bars represent the case of households' preferences becoming closer to those of the top income group.

Figure 5 CHANGES IN STOCKS AND FLOWS OF HOUSEHOLDS ACROSS QUALITY AND INCOME

7 Conclusion

We develop and estimate a structural model of housing demand and supply to quantify the distributional effects of leverage regulation. We match the demographic and financial characteristics of the household population in the capital of Norway, Oslo, to the universe of housing transactions between 2010 and 2018. Our model features housing decisions of financially constrained households and financially unconstrained investors and derives equilibrium prices via a market clearing condition. Our detailed data on income and wealth allows us to precisely measure households' affordability constraints due to Loan-to-Value (LTV) and Loan-to-Income (LTI) limits. We estimate households' moving costs and willingness to pay for neighborhood and property attributes across the income distribution.

We use the estimated model to conduct three counterfactual exercises. We first impose a tighter leverage limit relative to the baseline, which delivers the following results. First, we show that the housing choice sets of the lowest and highest income groups are only marginally affected by the change in LTI. This is so because low-income households already before the introduction of the LTI only had access to a small number of properties due to a binding LTV limit. For the top income group, the introduction of the LTI cap has little impact on their housing choice set. In contrast to these groups, middle-income households experience a reduction in their housing choice set of up to 18.6%, and up to 25% of homes in high-quality districts become out of reach. Second, for households that move, we show that a tighter LTI limit reduces the probability of becoming homeowners by 11.6% for low-income households but does not affect the richest. Overall, these findings highlight the inequality aspect of mortgage regulation.

Frictions that reduce real estate investors' incentive to trade in the housing market have two main effects. First, they increase house prices mostly for properties that are demanded by high-income households, who are the least price sensitive and have the highest willingness to pay from prime locations. As a result, such frictions have the indirect implication that they transfers wealth from those who move up the housing ladder to those that sit on top of the ladder. Second, they increase households' mobility, as they substitute for the reduction in investors' trading, and determine a reallocation of households from low- to high-quality districts, countervailing the regressive effect that a tighter LTI limit generates. Based on our findings, this happens because purchasing properties in low-quality districts, which are mostly held by investors, becomes more expensive, so households are better-off renting in higher quality neighborhoods.

Last, we simulate a scenario where we shift households' preferences closer to those of the top income group. This increases mobility, raises homeownership rates, and causes a

reallocation of households from low- to high-quality districts. Similarly to the previous counterfactual, this is another policy that regulators can implement to avoid the regressive effects of stricter LTI limits.

We leave several directions for future research. First, our methodological framework can be generalized by introducing an endogenous construction sector similarly to [Murphy \(2018\)](#). Such an extension would be useful to study the longer-term effects of leverage regulation on housing stock. We omitted this aspect from our application as we focus on a densely populated area where the supply of housing stock is restricted. Less densely populated geographic areas would represent a natural extension. Second, regarding the long-term effects of regulatory interventions, our framework could be extended with endogenous district quality, as in [Almagro and Domínguez-Iino \(2024\)](#). Changes in quality could be modeled as a function of changes in demographics in each area, which would make it possible to quantify the consequences of gentrification. Our current setup only has a single cross-section of district characteristics available. A richer dataset with variation in district quality over time would be a good starting point for such an analysis. Last, more detailed data on investors' portfolios and their asset allocation choice and affordability constraints would allow for more analysis of the interplay between households and professional investors in impacting social mobility and wealth inequality.

References

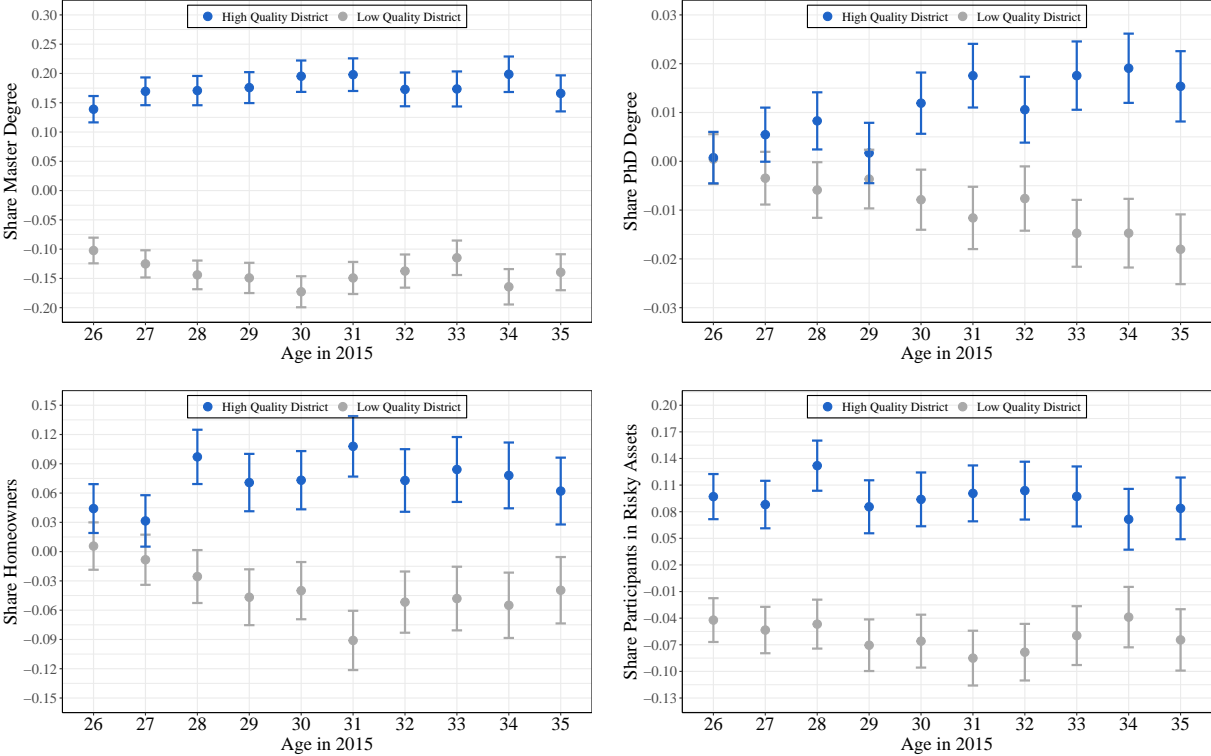
- Aastveit, K. A., R. E. Juelsrud, and E. G. Wold (2022). The leverage-liquidity trade-off of mortgage regulation. Norges Bank Working Paper.
- Acharya, V. V., K. Bergant, M. Crosignani, T. Eisert, and F. McCann (2022). The anatomy of the transmission of macroprudential policies. *Journal of Finance* 77(5), 2533–2575.
- Almagro, M. and T. Domínguez-Iino (2024). Location sorting and endogenous amenities: Evidence from amsterdam. Working Paper.
- Bach, L., L. E. Calvet, and P. Sodini (2020). Rich pickings? risk, return, and skill in household wealth. *American Economic Review* 110(9), 2703–2747.
- Bajari, P., P. Chan, D. Krueger, and D. Miller (2013). A dynamic model of housing demand: Estimation and policy implications. *International Economic Review* 54(2), 409–442.
- Baker, S. R. (2018). Debt and the response to household income shocks: Validation and application of linked financial account data. *Journal of Political Economy* 126(4), 1504–1557.

- Bayer, P., F. Ferreira, and R. McMillan (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy* 115(4), 588–638.
- Bayer, P., R. McMillan, A. Murphy, and C. Timmins (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica* 84(3), 893–942.
- Benetton, M., M. Kudlyak, and J. Mondragon (2022). Dynastic home equity. Working Paper.
- Brandsaas, E. E. and J. S. Kvaerner (2023). Partial homeownership: A quantitative analysis. Working Paper.
- Calvet, L. E., J. Y. Campbell, and P. Sodini (2009). Fight or flight? portfolio rebalancing by individual investors. *Quarterly Journal of Economics* 124(1), 301–348.
- Chetty, R. and N. Hendren (2018a). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *Quarterly Journal of Economics* 133(3), 1107–1162.
- Chetty, R. and N. Hendren (2018b). The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *Quarterly Journal of Economics* 133(3), 1163–1228.
- Chetty, R., N. Hendren, and L. F. Katz (2016). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review* 106(4), 855–902.
- DeFusco, A. A., S. Johnson, and J. Mondragon (2020). Regulating household leverage. *Review of Economic Studies* 87(2), 914–958.
- Epple, D., L. Quintero, and H. Sieg (2020). A new approach to estimating equilibrium models for metropolitan housing markets. *Journal of Political Economy* 128(3), 948–983.
- Fagereng, A., C. Gottlieb, and L. Guiso (2017). Asset market participation and portfolio choice over the life-cycle. *The Journal of Finance* 72(2), 705–750.
- Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh (2017). The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. *Journal of Political Economy* 125(1), 140–223.
- Geanakoplos, J. (2009). The leverage cycle. *NBER Macroeconomics Annual 2009* 24, 1–65.
- Han, L., C. Lutz, B. Sand, and D. Stacey (2021). The effects of a targeted financial constraint on the housing market. *Review of Financial Studies* 34(8), 3742–3788.

- Heckman, J. J. and R. Landersø (2022). Lessons for americans from denmark about inequality and social mobility. *Labour Economics* 77, 1–14.
- Holmes, T. J. and H. Sieg (2015). Structural estimation in urban economics. *Handbook of Regional and Urban Economics* 5A, 69–114.
- Kiyotaki, N., A. Michaelides, and K. Nikolov (2011). Winners and losers in housing markets. *Journal of Money, Credit and Banking* 43(2-3), 255–296.
- Koijen, R. S. J. and M. Yogo (2019). A demand system approach to asset pricing. *Journal of Political Economy* 127(4), 1475–1515.
- Ludwig, J., G. J. Duncan, L. A. Gennetian, L. F. Katz, R. C. Kessler, J. R. Kling, and L. Sanbonmatsu (2012). Neighborhood effects on the long-term well-being of low-income adults. *Science* 337(6101), 1505–1510.
- Mian, A. and A. Sufi (2009). The consequences of mortgage credit expansion: Evidence from the u.s. mortgage default crisis. *Quarterly Journal of Economics* 124(4), 1449–1496.
- Mian, A., A. Sufi, and E. Verner (2017). Household debt and business cycle worldwide. *Quarterly Journal of Economics* 132(4), 1755–1817.
- Murphy, A. (2018). A dynamic model of housing supply. *American Economic Journal: Economic Policy* 10(4), 243–267.
- Peng, Y. (2021). Mortgage credit and housing markets. Working Paper.
- Peydrò, J.-L., F. Rodriguez-Tous, J. Tripathy, and A. Uluc (2024). Macroprudential policy, mortgage cycles and distributional effects: Evidence from the uk. *Review of Financial Studies* 37(3), 727–760.
- Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica* 55(5), 999–1033.
- Sandlie, H. C. and J. Sørvoll (2017). Et velfungerende leiemarked? [“a well-functioning rental market?”]. *Tidsskrift for Valferdsforskning* 20(1), 45–59.
- Sommer, K., P. Sullivan, and R. Verbrugge (2013). The equilibrium effect of fundamentals on house prices and rents. *Journal of Monetary Economics* 60(7), 854–870.
- Stamsø, M. A. (2023). *Norway: Booming Housing Market and Increasing Small-Scale Landlordism*. Policy Press.

van Bakkum, S., M. Gabarro, R. M. Irani, and J.-L. Peydrò (2024). The real effects of borrower-based macroprudential policy: Evidence from administrative household-level data. forthcoming at *Journal of Monetary Economics*.

Appendix A Figures



Notes: The figure plots the regression coefficients $\hat{\eta}_j$ from equation (1) with confidence intervals.

Figure 6 LIFE OUTCOMES WHEN GROWING UP IN LOW VS HIGH-QUALITY NEIGHBORHOOD

Appendix B Data

This section provides additional details on how we construct our sample.

B.1 Variable Definitions

For each individual in our sample, we observe the birth date (variable name: “foedsels_aar_mnd”) from the population database (In Norwegian: “Befolkning”). In the same database, we observe the number of child(ren) each individual has and the birth dates of the child(ren) (variable name: ‘fodselsdato_barn_01-10”). In addition, we observe the ID (anonymized) of the spouse (variable name: ‘ekt_fnr_aaaa”), or cohabitant (variable name: “sambo_snr_aaaa”). We use this information to classify an individual into a one-adult household (not registered ID for spouse or cohabitant) or more than one adult household. We refer to the two household types simply as singles and couples. The tax authority collects information on the

complete wealth holdings of all households at the end of every year. For tax purposes, the household can allocate wealth in a way that gives the lowest wealth tax. Thus, there are no incentives for tax-motivated asset allocation within the household.

The financial information comes from the Norwegian Tax Registry (NTR) and reflects individuals' tax returns. We obtain this data from Statistics Norway, which merges it with the above demographic data. The NTR is responsible for collecting income and wealth taxes in Norway. By law, employers, banks, and public agencies must disclose personal information on income and wealth to the Tax Administration. The tax return includes all sources of income, as well as detailed information on wealth and debt. Individuals are accountable for the information provided in their tax returns, and the submission of inaccurate information is punishable by Norwegian law.

For each individual, we include total income (variable name: "wsaminnt"), which is the sum of gross salary income and pension plus net capital income and total government transfers, debt (variable name: "gjeld"), the value assessment of principal residence, which we label simple real estate (variable name: "prim_mark"), financial wealth (variable name: "bruttofin"), and total assets (variable name: "ber_brform"). We define homeownership as a variable that takes the value of one for all individuals with a positive value assessment of the principal residence. We also calculate adjusted total assets as total assets minus real estate. As we explain below, we net out the value assessment of the home because we include it in the estimation of the household's net worth, which is based on house product prices, $P_{j,t}$.

With this data, we need to aggregate individual data into household data. To be included in our sample, we follow standard practice in household finance and ensure that the households included in the analysis have a minimum cash balance (see, e.g., [Calvet et al., 2009](#); [Fagereng et al., 2017](#)). In our case, we require the household to have at least 5,000 NOK in financial wealth. Because we study the mobility pattern of households, we also exclude from our sample the 20% with the lowest income at the end of the year and restrict the sample to households that are at least 18. Everyone that satisfies these criteria and does not have a registered spouse or cohabitant is classified as single.

The corresponding definition of couples is a bit more involved. We start by calculating the same financial data as for singles. We then aggregate total income, adjusted total assets, debt, and financial wealth to the household level. For age and homeownership we select the maximum in the household. We keep the ID (anonymized) of the oldest individual in the household and refer to this individual as the household head. District $d \in \mathcal{D}$ of residence and the number of children in the household are based on the household head. For the transactions, we include all transactions done by any of the two adults in the household. We impose the same financial and age requirements on couples as singles. The sum of single

and couples satisfying our basic requirements comprise our sample of households. For each household, we define net worth, $A_{i,t}$, as:

$$A_{i,t} = P_{j,t} + \text{Financial Wealth}_{i,t} + \text{Other Real Estate}_{i,t} - \text{Debt}_{i,t}, \quad (27)$$

where $P_{j,t}$ is the price for house product $j \in \mathcal{J}$ in year t .

B.2 The Sample

We need the number of rooms and district for each unit and information about the homeowner or the renter to estimate the model. If these data are missing, we impute them. In our final sample, all homeowners and renters live in a particular housing product. In addition to households, we define a real estate investors' sector that owns part of the housing stock. In what follows, we explain how we deal with missing values and define homeowners, renters, and investors.

Regarding housing characteristics, the number of rooms $u \in \mathcal{U}$ is missing for 14,801 transactions (4.4%). For those observations, we use a multinomial logistic regression to predict the number of rooms based on the size of the apartment, the transaction price, and the district. The model predicts correctly in 74% of the cases. In comparison, randomly selecting the number of rooms $u \in \mathcal{U}$ would predict correctly in only 20% of the cases.

Regarding information about the homeowner or the renter, we observe it for households who transact in the market. For everyone else, we predict their housing product. Because we know the district where each household lives every year, it is sufficient to predict the number of rooms $u \in \mathcal{U}$ to identify their house product $j \in \mathcal{J}$. We begin by selecting all transactions in our sample period in which a household purchases a house with the number of rooms $u \in \mathcal{U}$. We use a rich set of characteristics for this sample to predict the number of rooms in their units. These characteristics include age, age², age³, a dummy variable for being single, number of children, total income, and financial wealth. And the following dummies: D_{1i} takes the value of 1 if household i 's total income is in the top 10 percent of the income distribution, D_{2i} takes the value of 1 if household i 's financial wealth is in the top 20 percent of the financial wealth distribution, and D_{3i} takes the value of 1 if household i has more than four children. The idea with the indicator variables is to let income and wealth matter differently for very wealthy individuals relative to the rest of the sample. The model predicts correctly in approximately 54% of the cases. In comparison, randomly allocating the number of rooms $u \in \mathcal{U}$ would give a success rate of 33.3%.

Having identified a model that takes household characteristics as input and assigns the number of rooms $u \in \mathcal{U}$ as output, we create our sample. To do so, we start by selecting all

households who live in any district $d \in \mathcal{D}$ at the end of 2010. For those households that also bought a housing product in the same year, we assign their actual product choice j to them. For the remaining households, regardless of homeownership status, we predict their housing product $j \in \mathcal{J}$ as we now explain.

Starting with 2010, the first year in our sample, we predict the number of rooms $u \in \mathcal{U}$ in the housing unit for all households for which we do not observe it. Given that we use a multinomial logistic regression model for this prediction, the output is a probability distribution for the number of rooms $u \in \mathcal{U}$. Since we have data on the number of housing units with $u \in \mathcal{U}$ number of rooms in each district $d \in \mathcal{D}$, we ensure that we never assign more housing units to a particular type $u \in \mathcal{U}$ than what is reported in official statistics. In addition, we ensure that the relative frequency distribution of housing units with $u \in \mathcal{U}$ number of rooms match official statistics. Given these restrictions, we assign the most likely choice, as predicted by our model, to each household.

An example illustrates what we do. Assume official statistics report that in district $d = 1$ there are 1,000 units with one room ($u = 1$) and 2,000 units with two rooms ($u = 2$). The total number of households in district $d = 1$ is 10,000. In our sample, assume that 8,000 households that live in district $d = 1$ satisfy the requirements to be included in the sample. Of those 8,000, we observe 500 households buying a housing unit with one room ($u = 1$) and 500 buying a housing unit with two rooms ($u = 2$). The number of households for which we need to predict the housing product is $8,000 - 1,000 = 7,000$. We then assign the housing product $j = \{(1, 1) | d \in \mathcal{D}, u \in \mathcal{U}\}$ to: $\max\{(\text{Total } j \text{ units} / \text{Total households in } d) \times \text{Total households in } d \text{ in our sample} - \text{Number of housing we observed buying product } j, 0\}$, which in this example is $\max\{1,000/10,000 \times 8,000 - 500, 0\}$.

In the next step, we first exclude the 800 households we just assigned a housing unit, then repeat the exercise for housing units with two rooms ($u = 2$). We continue until all the households in district $d = 1$ have a housing product. In all other years (i.e., the period from 2011 to 2018), we use the same method to assign the number of rooms $u \in \mathcal{U}$ in a housing unit for households that enter the sample without buying a housing unit or move to another district. Households entering the sample by purchasing a housing unit are given the housing product $j \in \mathcal{J}$ they choose.

To separate homeowners from renters, we calculate for each year the average book value of housing for those that transacted in the market during the year. Households that do not transact in the market but have a book value of housing above this estimate are classified as homeowners, while the rest of the households are renters. The threshold for homeownership is product specific.

In addition to the household sector, we include a real estate investors' sector that trans-

acts in the housing market to maximize risk-adjusted profits. We define investors as all market participants that bought or sold a housing unit $j \in \mathcal{J}$ from 2011 to 2018 that do not satisfy our requirements to be classified as households. In addition, we classify households that buy or sell multiple units in a year as investors.

B.3 Determinants of Mortgage Interest Rates

In this section we provide evidence to support our modeling assumption, driven by data limitations, that a change in LTI limits has little to none effect on mortgage interest rates, while primarily affecting households' choice sets instead. We have access to mortgage interest rates' data for the mortgage products offered by the five largest Norwegian mortgage providers, with combined market share in 2020 of around 70% in Oslo, for the period between 2008 and 2018.²⁶ In the regressions displayed in Table 9 we show that 81 percent of the variation in mortgage rates is explained by the Norwegian Central Bank's policy rate and by dummies for mortgage product characteristics, while bank fixed effects only explain an extra 1% of interest rate variation. All else equal, we also show that after the introduction of the LTI limit, captured by a dummy variable, there was no statistically significant change in mortgage rates. These results, together with the evidence from Figure 2 on house price effects of LTI changes, leads to believe that the first order channel through which LTI limits affect housing markets is via changes in households' choice sets rather than changes in mortgage interest rates.

More specifically, in the first column of Table 9 we show that the mortgage interest rate markup over the policy rate is about 3.1%, and that the pass-through of the policy rate is 0.7. Alone, the policy rate explains 33% of the variation in mortgage rates. The second column adds second and third-order polynomials of the policy rate, which explain an extra 1% of the variation. This specification has more explanatory power than one with only month-year fixed effects, reported in the third column. In column four we adds dummies for mortgage product characteristics, which brings up the R-squared to 81%. In the fifth column we include bank dummies, which only increase the R-squared by 1%, and in the last column we add a post-LTI dummy, which is not statistically significant.

²⁶We downloaded these mortgage data from finansportalen.no.

Table 9 DETERMINANTS OF MORTGAGE INTEREST RATES

	Nominal Mortgage Interest Rate (Including Fees)					
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	3.12*** (0.14)	2.60*** (0.27)	4.14*** (0.16)	2.71*** (0.34)	2.76*** (0.29)	2.77*** (0.28)
Post-LTI						-0.01 (0.08)
Policy Rate	0.70*** (0.04)	1.51** (0.34)		1.48** (0.34)	1.47** (0.33)	1.46** (0.33)
(Policy Rate) ²		-0.31** (0.08)		-0.35** (0.08)	-0.35** (0.08)	-0.35** (0.09)
(Policy Rate) ³		0.03** (0.01)		0.04*** (0.00)	0.04*** (0.)	0.04*** (0.01)
DNB					-0.21*** (0.03)	-0.21*** (0.03)
Handelsbanken					-0.38** (0.13)	-0.38** (0.12)
Sparebanken					0.22*** (0.03)	0.22*** (0.03)
Nordea					-0.01 (0.03)	-0.01 (0.03)
Month-Year FE	No	No	Yes	No	No	No
Mortgage Product FE	No	No	No	Yes	Yes	Yes
Observations	1,940	1,940	1,940	1,940	1,940	1,940
Adjusted R ²	0.33	0.34	0.33	0.81	0.82	0.82

Notes: This table reports the results of six regressions where the dependent variable is the monthly nominal mortgage interest rate (that include origination fees) between January 2008 and June 2019.. The data includes 8 mortgage products originated by the 5 main banks (with combined market share in the Norwegian mortgage market of around 70%). Mortgage products are defined by fixed vs variable interest rate, 5 LTV buckets ($\leq 60\%$, 60-70%, 70-75%, 75-80%, 80-85%), and three fixation periods for fixed rates mortgages (3, 5, 10 years). The regressions include the Norwegian central bank policy rate, including a second and third-order polynomial, mortgage product fixed effects, month-year fixed effects, and a dummy variable that takes the value of one after the introduction of the LTI limit in January 2017. Standard errors are clustered at the bank level.

Appendix C Tables

Table 10 LONG-TERM (TWO YEARS) EFFECT OF CHANGE IN LTI LIMIT ON CHOICE SETS AND MOVING PROBABILITIES

	LTI	Income Groups				
		1	2	3	4	5
All households						
Owners' Share of Total Products in Choice Set	Base	52.9%	61.3%	68.4%	73.3%	73.1%
	Δ	-14.3%	-18.8%	-17.5%	-15.2%	-1.0%
Owners' Share of High Quality Products in Choice Set	Base	14.2%	16.0%	18.7%	21.3%	22.9%
	Δ	-12.5%	-19.4%	-26.2%	-20.8%	-3.2%
Owners' Moving Probability	Base	4.3%	4.0%	3.9%	4.2%	5.9%
	Δ	-0.3%	-0.6%	-0.7%	-0.9%	-0.1%
Renters' Moving Probability	Base	16.5%	13.8%	11.5%	11.1%	13.5%
	Δ	-0.8%	-2.2%	-1.8%	-1.5%	-0.1%
Movers						
From Renting to Owning	Base	5.5%	15.2%	27.7%	39.0%	44.1%
	Δ	-12.2%	-11.4%	-4.3%	-3.2%	-1.1%
From Low to High Quality	Base	10.4%	13.2%	15.1%	17.7%	16.5%
	Δ	-1.2%	-0.7%	-1.6%	-1.6%	-0.2%

Notes: This table reports the longer-term effect of changes in LTI limits on the mobility of households for each income quintile. The baseline scenario (Base) corresponds the case without changes in LTI limits. The counterfactual scenario of changing LTI limits has been introduced in the previous period, and the percentage changes under counterfactual scenario of changing LTI limits are reported (Δ). The first panel reports the share of available products for home purchasing (Owners' Share of Total Products in Choice Set), the share of available high-quality products for home purchasing (Owners' Share of High-Quality Products in Choice Sets), the probability of moving for homeowners (Owners' Moving Probability), and the probability of moving for renters (Renters' Moving Probability). The second panel considers only movers, where the probability of movers changing from renting to owning (From Renting to Owning) and the probability of movers changing from low-quality districts to high-quality districts (From Low to High Quality) are reported.