Leverage Regulation and Housing Inequality

Jens Soerlie Kvaerner, Nicola Pavanini, Yushi Peng*

July 2025

Abstract

We estimate an equilibrium model of housing demand and supply, quantifying the distributional effects of leverage regulation on household mobility, access to high-quality housing, debt and house prices. 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. We recover households' willingness to pay for housing quality and moving costs. Our counterfactuals quantify costs and benefits of loan-to-income (LTI) limits. While tighter limits reduce household debt and house prices, they also have regressive effects on mobility. We document how these effects depend on household preferences and financial constraints, and can be offset with housing subsidies, increasing housing supply, or with rent controls.

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

*Kvaerner: BI Norwegian Business School, jkverner@gmail.com. Pavanini: Tilburg University and CEPR, n.pavanini@tilburguniversity.edu. Peng: Tilburg University, y.peng@tilburguniversity.edu. We benefitted from the feedback of William Burton (discussant), José Ignacio Cuesta, Andreas Fagereng, Andreas Fuster, Alessandro Gavazza (discussant), Zhonglin Li (discussant), Clara Martnez-Toledano (discussant), Victoria Vanasco (discussant), Stijn Van Nieuwerburgh, Yu Zheng (discussant), and the seminar and conference participants at Copenhagen Business School, Tilburg, Bank of England, CEPR Endless Summer Conference, PUC Chile, EIEF-UniBo-IGIER IO Workshop, Bank of Italy, CEPR Household Finance Conference, Bristol IO Workshop, UZH Finance, IIOC, UNIL-EPFL, FIRS, CEPR Applied IO Conference, SHUFE IO, EARIE, SITE Financial Regulation, Erasmus University Rotterdam, ECB, Boston College, Bank of Greece, Norges Bank, Central Bank of Ireland, LSE, CEMFI, Banque de France, IWH Halle, USI Lugano, ECARES, Deakin, Melbourne. 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, for most households, real estate is both the largest investment they make and the only one they finance with debt. Because of leverage, even small house price changes can lead to large gains or losses in wealth. As a result, housing is a key determinant of where household end up in the wealth distribution at retirement age, and of wealth inequality more broadly (Fagereng et al., 2021; Bach et al., 2020).

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.

Recent evidence documents 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). However, we know little about the distributional effects of these policies on home-ownership and household mobility, and the extent to which the regressive effects of these interventions can be offset with policies such as housing subsidies, expansion of the housing stock, or rent controls. Uncovering these effects requires a model that quantifies how leverage limits and other housing policies impact house prices, home-ownership and housing mobility differentially across households' demographics, separately identifying the role of households' financial constraints from preferences for neighborhood quality in residential choices.

Our paper addresses this task developing a structural model of housing demand, supply, and equilibrium pricing 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 that document benefits and costs of tighter LTI limits, comparing them to the effects of progressive housing policies. On the positive side, we quantify how stricter limits

decrease household debt and house prices. On the negative side, we show how such policies reduce mobility choices and access to high-quality housing in a regressive way. 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. More generally, we document how macro-prudential policies shape housing segmentation and redistribute wealth persistently across households, through general equilibrium housing channels.

Our framework extends the most recent structural models of residential choices, mostly from the industrial organization and urban economics literature. 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 their residential property and face transaction costs, and real estate 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, similar to the one used by Fagereng et al. (2021). Among its many advantages, these data provide particularly detailed information on households' income and wealth, necessary to investigate distributional implications of leverage and housing policies.[1] 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

[1]While Norway has lower income inequality than the U.S., the two countries have comparable wealth inequality, as in 1993-2000 the average top 0.1% wealth share was 12.5% in Norway and 13.6% in the U.S. (Fagereng et al., 2020). According to the World Inequality Database (`https://wid.world/`), in 2010-2019 Norway's average income (wealth) inequality, defined as the share of income (wealth) held by the top 1%, was 10.9% (22.7%), similar to France with 10.6% (25.3%), Germany with 12.7% (27.5%), and the United Kingdom with 12.6% (21.3%).

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.

We document mobility patterns across households' income, district quality, and homeownership status. This allows us to show the main sources of variation that identify our structural parameters, and to justify our modeling choices. We display yearly transition matrices that report households' likelihood of remaining in their current property or moving to a different type of residence, as a function of their income, quality of their current district, and homeownership status. These descriptives show that housing choices are sticky, as the probability of remaining in the same housing product is on average around 90%. However, conditional on moving, there is significant heterogeneity in mobility patterns. While low-income households are more likely to transition within rented properties and have higher likelihood of downgrading, the opposite is true for high-income households, who transition mostly within the ownership market and have higher likelihood of upgrading.

To provide descriptive evidence on the importance of leverage regulation, we show the impact of introducing an LTI limit on households' choice sets and leverage. In 2017, Norway set the LTI limit to 5. Before that, there was no strict LTI threshold. We divide housing products across size and quality, and document heterogeneity in how the LTI limit affects the share of affordable housing products across income levels. We also show that home buyers in the middle of the income distribution exhibit bunching in their Debt-To-Income around the threshold right after the introduction of the limit, consistent with the new limit representing a binding financial constraint, particularly for low and middle-income households.

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, they incur moving costs, expressed 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. Data on homeownership allows us to distinguish between five alternative decisions households can make, which until now have not been fully captured in the literature. 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 moving costs across the income distribution. Higher income households face lower costs, while older households and bigger families face the larger costs. Second, willingness to pay for housing quality and homeownership increases with income. Our model predictions fit well several empirical moments, with a close match of households' mobility, for various demographics, across rental and ownership and across district quality.

We simulate counterfactual scenarios to quantify the distributional effects of leverage limits, and to investigate whether three commonly used housing policies can mitigate the regressive effects of these limits. First, we introduce a tighter LTI limit than what was implemented in 2017, setting it to 3 instead of 5. We document, as expected, that a more stringent LTI limit reduces the share of affordable housing products differently across the income distribution. The choice sets of the lowest and highest income groups are largely unaffected. For lowest-income households, this is because the former LTI limit already restricted their choice set to a small fraction of properties. For highest-income households, lowering the LTI limit from 5 to 3 is 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 18.6% in the share of affordable housing products. We then zoom in on households who moved during our sample period. 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 up to 43.5% for low-income households, 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.

This counterfactual also provides evidence on the benefits of tighter LTI limits. First, we document that aggregate household leverage declines across all income groups for households purchasing a new home, with the larger drop of up to 34% experienced by those in the middle of the income distribution. Second, we find that a drop in the LTI limit from 5 to 3 causes a reduction in house prices ranging between 2% and 5.1%. While this price reduction is beneficial especially for first time buyers, it also implies a decline in wealth for homeowners.[2]

Three takeaways emerge from this counterfactual. First, we document that tight leverage limits have distributional implications along two dimensions. On the one hand, they reduce homeownership disproportionally more for low-income households. On the other hand, they shrink access to high-quality districts disproportionally more for middle-income households, restricting their housing choice set and probability of upgrading along the housing ladder. Second, we quantify the extent to which segmentation in housing markets is driven by financial constraints. Even without stricter LTI limits, high-income households are already more concentrated in high-quality neighborhoods as owners, indicating existing segmentation potentially driven by preferences or financial constraints. Introducing stringent LTI limits reinforces this pattern and further exacerbates housing segmentation, highlighting the role of financial frictions. This justifies the introduction of progressive housing policies simulated in our following counterfactuals. Third, a strict LTI threshold reduces house prices the most for low-quality properties, as low-income households become less likely of being homeowners.

We conduct three additional counterfactuals to simulate housing policies that have the potential to offset these regressive effects of LTI limits. First, we introduce a progressive housing subsidy, and document that it encourages mobility, especially for low-income households, and creates an incentive to move to high-quality neighborhoods. Second, we expand housing supply in high-quality districts, and show that this has a similar effect as the subsidy, as it favors reallocation of households from low and middle-quality neighborhoods to high-quality ones. Last, we simulate a cap on rental prices that is predominantly binding for low-quality districts. This policy encourages reallocation of low-income households from high

[2]Another benefit of tight LTI limits is lower household default risk. We are not able to directly quantify this effect due to lack of data on mortgage performance, but the reduction in household debt that we document is likely a good proxy for lower defaults. A reduction in credit risk is however unlikely to be the primary benefit that regulators in Norway target, as mortgage default rates of Norwegian households have been particularly low at around 1% since the 1990s (Solheim and Vatne, 2013), with losses on mortgage loans for banks averaging 0.2% (https://www.norges-bank.no/en/topics/Statistics/bankstatistikk/). The reason for these low default rates and losses is that upon default the bank forces house liquidation, and the mortgages are full recurse, with lenders having the right to garnish wages and social security transfers.

and middle-quality areas to low-quality ones, but also favors low-income households' homeownership in low-quality districts, as investors find it less profitable to own those properties when rents are controlled.

Related Literature. We contribute to three strands of the literature. The first consists of structural equilibrium housing models. These frameworks are based on three approaches. First, dynamic discrete choice models from industrial organization and urban economics. Examples are Bayer et al. (2016) and Epple et al. (2020), while Holmes and Sieg (2015) provides a comprehensive survey of the literature. This is the methodology we employ for our analysis. Closest to our paper are Peng (2023), which incorporates financing conditions into housing decisions with search frictions, but without household-level financial data, and Almagro and Domínguez-Iino (2024), who build a model of residential sorting that quantifies the importance of endogenous location amenities for inequality. Ouazad and Rancière (2018) develop a general equilibrium model of housing choice with borrowing constraints, estimate a static housing demand model, and focus on mortgage approval. Relative to this literature, our contribution is to incorporate affordability constraints linked to leverage limits, to model dynamic housing decisions of both households and investors in ownership and rental markets, to allow for vertical heterogeneity in housing products, and to model equilibrium prices.[3]

The second approach is of life-cycle models from the household finance literature (Cocco, 2005; Gomes and Michaelides, 2005; Attanasio et al., 2012; Bajari et al., 2013; Kaplan et al., 2020). The dynamic discrete choice approach shares several similarities with life-cycle models. In both cases, housing choices result from dynamic utility maximization with transaction costs, where households are heterogeneous and form expectations over some relevant state variables when making decisions. When it comes to the differences, the life-cycle approach mostly relies on calibration or estimation based on aggregate moments, is primarily focussed on households' heterogeneous decisions across the life-cycle, and allows for optimal consumption and savings choices. However, it typically considers a single type of housing product, or allows at most for limited heterogeneity in housing size and location.

On the other hand, the dynamic discrete choice approach relies on estimating all parameters of households' utility and transaction cost functions based on individual households' decisions, allows for household heterogeneity across income, wealth, and demographics, as well as extensive heterogeneity in utility from housing property attributes such as size and location. It does not explicitly consider consumption and saving choices, but is well suited to model equilibrium house prices via market clearing for a wide range of differentiated hous-

[3]A similar discrete choice approach has been used also to model mortgage demand and supply (Benetton, 2021; Buchak et al., 2024; Benetton et al., 2025). We complement that literature modeling housing demand and supply, and allowing households' housing choice set to depend on their maximum mortgage size.

ing products.[4] We rely on the dynamic discrete choice approach due to our interest on the distributional effects of leverage limits and housing policies across households' demographics and heterogeneous housing products, with emphasis on ownership, mobility, and equilibrium prices. Closest to our work are spatial equilibrium life-cycle models that account for location choices (Favilukis et al., 2023; Mabille, 2023; Fonseca et al., 2024).

The third approach is that of 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; Piazzesi and Schneider, 2016). A common limitation of these papers is that they consider households' financing decision mainly from an aggregate perspective, preventing them from analyzing distributional implications with the same degree of heterogeneity as our model. We do not adopt this approach because our focus is not on macroeconomic general equilibrium implications of leverage regulations, but rather on their distributional effects on housing markets.

A second strand of the literature we contribute to is on empirical evidence of the effects of macro-prudential regulations. 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 Bekkum 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. One exception is Tzur-Ilan (2023), who documents with a difference-in-differences how the introduction of LTV limits in Israel changed borrowers' location decision. We differ by using a structural model that allows us to simulate equilibrium responses to

[4]We show with hedonic pricing regressions that the degree of product aggregation implied by standard life-cycle models explains around 22% of the variation in house transaction prices, while the heterogeneity in our framework explains 66% of the price variation.

leverage limits and housing policies, and by modeling housing choices for the population of households rather than just mortgage borrowers. 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, or how other housing policies could offset the regressive effects of leverage regulation. 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 and housing policies. This is the framework that we develop in this paper.

Last, we contribute to the literature on the effects of social mobility on households' economic outcomes and inequality. 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.

2 Data and Stylized Facts

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

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 res-

idential real estate market, both for individual properties and for portfolios of properties.[5] 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.[6] We explain all district-level variables and how we use them in Section 2.2.1.

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.[7] 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 $h \in \mathcal{H} = \{0, 1\}$ that takes the value of one for all individuals with a positive value assessment of the principal residence, and zero for renters.

We distinguish between individuals living alone and 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 as singles and couples and let G denote the set of demographic variables, including family size and age of the household head. 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. 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

[5] More information is available here: https://eiendomsverdi.no/

[6] Data is freely available at https://statistikkbanken.oslo.kommune.no/webview/

[7] 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 submitting inaccurate information is punishable by law.

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.

2.1.1 The Housing and Mortgage Market in Norway

We now provide the essential institutional details about the Norwegian housing and mortgage market. 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.[8]

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.[9] In December 2016, the FSA introduced an LTI limit equal to five times the gross annual household income. Because the frequency of our household data is annual, we assume that the LTI limit was introduced in 2017.[10]

The Norwegian mortgage market is dominated by two large banks, DNB and Nordea.[11] Most mortgages feature variable interest rates with a 20-30-years maturity. Fixed-rate mortgages 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

[8] For more information about the Norwegian housing market, we refer to Brandsaas and Kvaerner (2023), Sandlie and Sørvoll (2017), and Stamsø (2023).

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

[10] Many households obtained pre-qualification letters in the months before the introduction of the limit, hence the LTI constraint did not bind for everyone before early 2017.

[11] Their 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`.

do not choose which mortgage product to get to finance a property, but rather face a yearly user cost of housing, approximated with rental costs, 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 C.1, 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 $\mathcal{D} = \{1, 2, ..., 15\}$. 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 $\mathcal{J} = \{1, 2, ..., 45\}$ housing products.

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 A1 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 price across all transactions for product j.

We benchmark the degree of price variation that our model allows for to other standards in the literature, namely the macroeconomic and the life-cycle approaches to housing demand. Table A2 summarizes the results from five regressions where the log of house prices is the dependent variable, where we display the explanatory power of various dimensions of house price heterogeneity. The first specification includes only time-series variation - month and year fixed effects - and accounts for 22% of the variation in house prices. This is a standard benchmark in macroeconomic housing models that abstract from micro-level heterogeneity. As subsequent columns demonstrate, incorporating richer housing size and location-level

fixed effects enables us to retain a greater proportion of the total variation in house prices. The final specification, which includes fixed effects for zip codes (406 categories) and controls for house size in square meters, explains 80% of the variation in log prices, capturing both temporal and granular spatial differences in pricing.

The third specification is the one we use in our structural model. It controls for room-by-district fixed effects (45 combinations), in addition to month and year fixed effects, and achieves an R^2 of 0.66. This strikes a balance between capturing rich heterogeneity and maintaining computational tractability in a model setting. Importantly, it retains three times more explanatory power than the baseline time-series-only specification and approximately 1.5 times more than a common approach in the life-cycle literature that categorizes houses into three broad-size bins. This suggests that relatively simple spatial-product interactions go a long way toward explaining house price variation and that omitting such heterogeneity may result in a substantial loss of information about the price structure.

2.2 Sample and Summary Statistics

We construct our sample dynamically. We begin selecting all households who live in Oslo at the end of 2010. We follow the same approach for the following years, and assume that a household moves out of the city in year t if it was living in Oslo in year $t-1$ but is not there anymore in year t. 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 3.3 million observations. To provide a snapshot of the data, Table 1 presents descriptive statistics for the year 2015.

2.2.1 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 secondary school, the number of reports to the child welfare service per capita,[12] 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. The second survey

[12]Total number of reports by district-year is divided by district population and multiplied by 1,000. This normalization is applied to make all the variables in comparable units.

Table 1 Descriptive Statistics

	Homeowners				
	Mean	Std Dev	10th	50th	90th
Fraction couples	0.75	0.43	0.00	1.00	1.00
Age	48	16	30	46	71
Number of children	1.3	1.2	0.0	1.0	3.0
Total income	751	1,937	303	558	1,140
Gross wealth	5,209	26,323	2,496	3,567	7,409
Debt	1,518	2,355	9	1,128	3,328
	Renters				
	Mean	Std Dev	10th	50th	90th
Fraction couples	0.54	0.50	0.00	1.00	1.00
Age	41	15	26	37	65
Number of children	0.91	1.27	0.00	0.00	3.00
Total income	479	949	206	400	745
Gross wealth	1,025	15,835	7	112	2,761
Debt	544	1,624	0	113	1,734

	Homeowners		**Renters**	
	N of Obs	Share	N of Obs	Share
Stayers	189,309	0.88	84,344	0.72
Entering Oslo	5,486	0.03	9,382	0.08
Move within district in Oslo	8,093	0.04	7,942	0.07
Move between district in Oslo	12,836	0.06	15,828	0.13
Total	215,724	1.00	117,496	1.00

Notes: This table reports descriptive statistics of our sample in 2015. The top two panels report demographic and financial data, the bottom panel reports mobility statistics. We report the descriptive statistics for homeowners and renters separately. Besides the reported households living in Oslo in 2015, there are 12,702 owners and 13,493 renters who left Oslo that year. All financial variables are reported in NOK thousands.

question is the score on a happiness index.[13] We use the scores at the end of 2015 because it is the only year we observe scores on all five quality indicators.

Table A3 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. The first column reveals large price differences for

[13]The two questions are respectively: "How satisfied or dissatisfied are you with your health?" and "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.

a 3-room apartment between neighborhoods. The high-quality neighborhood is about 30% more expensive than the baseline, 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.

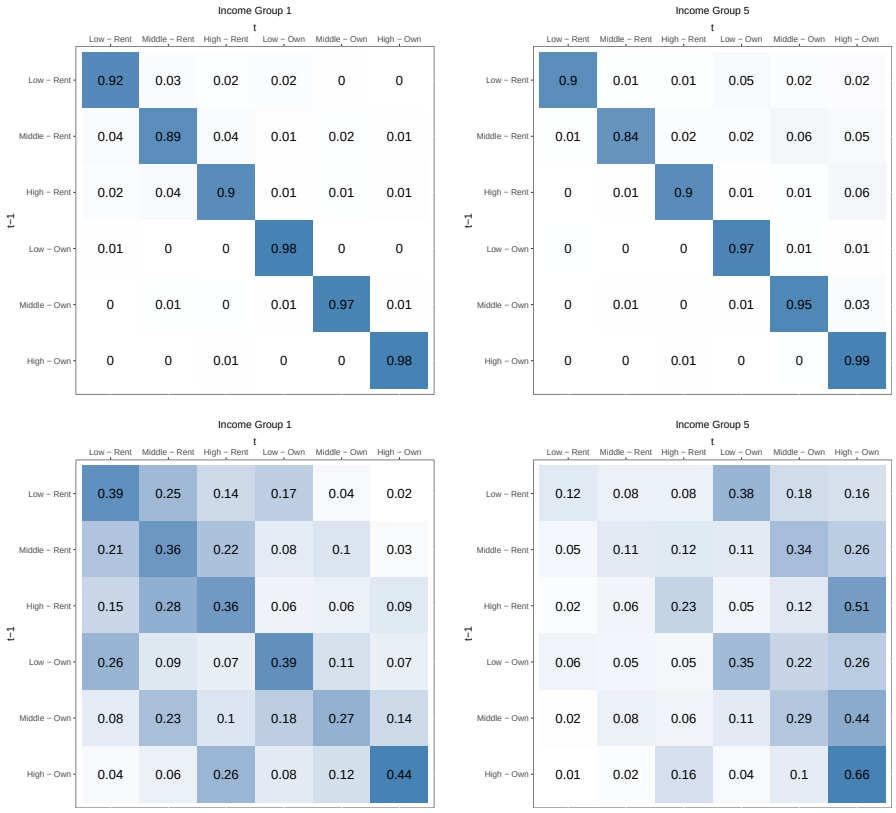
2.3 Mobility Patterns

One of the main predictions that our model delivers consists of households' probability of moving across housing products over time, allowing for differences in mobility patterns across their income, wealth, and demographics, and across housing quality, size, and homeownership status. Figure 1 presents four transition matrices to justify the choice of these dimensions of heterogeneity, and to display descriptive evidence of mobility patters within Oslo that drive the identification of our structural parameters on preferences and moving costs.

Each cell within each matrix shows the probability that in year t a household will live in a rented property (first three columns) or own a house (last three columns) in a low/middle/high quality district, depending on where the household was in year $t-1$ among those same six alternatives, displayed on the left vertical axis. The top two figures show transition matrices for all households, while the bottom two for movers. The two left figures show transition matrices for the lowest (1^{st}) income group, while the two right ones for the highest (5^{th}) income group. The upper diagonal elements of each matrix can be interpreted as probabilities of upgrading, that is going from renting to ownership or moving to a district of higher quality, while the lower diagonal elements represent probabilities of downgrading.

These figures convey two key facts that justify our modeling choices. First, as the top two figures show, housing choices are sticky, as the probability of remaining in the same housing product is on average around 90%.[14] This highlights the importance of modeling fixed moving costs and to include such large fraction of households who hardly ever move to model equilibrium prices. Second, as the two bottom figures show, there is significant heterogeneity in mobility patterns across income, housing products, and homeownership status. While low-income households are more likely to transition within rented properties and have higher likelihood of downgrading, as moving probabilities are higher in the lower diagonal elements, the opposite is true for high-income households, who transition mostly within the ownership market and have higher likelihood of upgrading.

[14]The transition matrices report household mobility patterns within Oslo, while around 4% to 8% of homeowners and 8% to 14% of renters leave Oslo each year.



Notes: These figures plot probabilities of moving, averaged across time within our sample period. The left vertical axis refers to locations and ownership of households in year $t-1$, while the top horizontal axis refers to locations and ownership of households in the following year t. Locations can be of three quality levels ("Low", "Middle", "High"), and within these locations households can be renters ("Rent") or homeowners ("Own"). For example, the top-left cell of each figure is the probability that a household who rents in a low-quality district in year $t-1$ remains in a low-quality district as renter in year t. The two top figures represent all households, while the bottom two figures represents only movers. The two left figures represent the lowest (1^{st}) income group, while the two right figures represent the highest (5^{th}) income group. For ease of interpretation, these are all probabilities conditional on not moving out of Oslo. The middle income groups' transition matrices are not included as their probabilities change roughly linearly moving from the lowest to the highest income group, so showing the lower and upper bounds is sufficient.

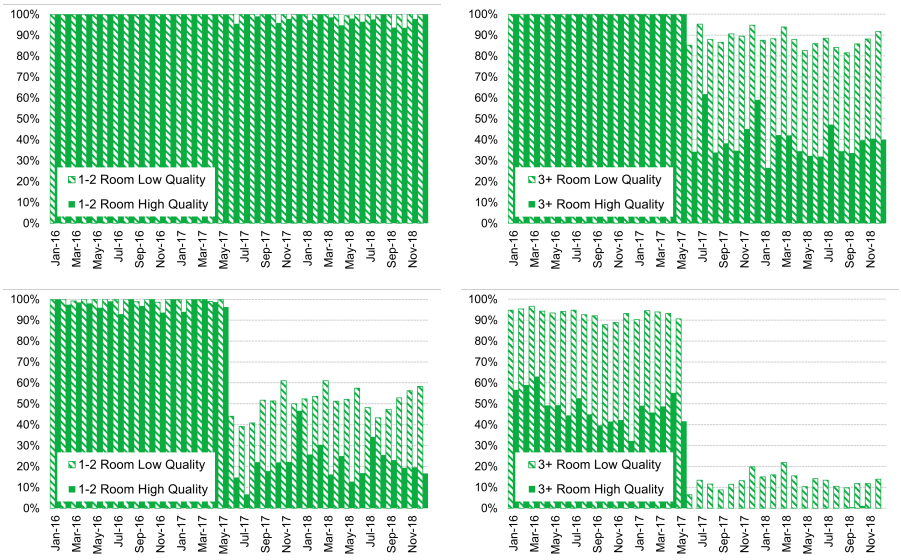
Figure 1 Mobility Across Homeownership, Income, Housing Quality

2.4 Impact of LTI on Prices, Choice Sets, Leverage

Our last preliminary analysis shows the effect of the loan-to-income (LTI) cap on house prices, households' choice sets, and households' leverage. The event we study is the introduction of an LTI cap in Norway in 2017. Figure A1 suggests that the LTI regulation 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 and four housing categories, 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 with Figure 2 how the LTI regulation affected individuals' choice sets.

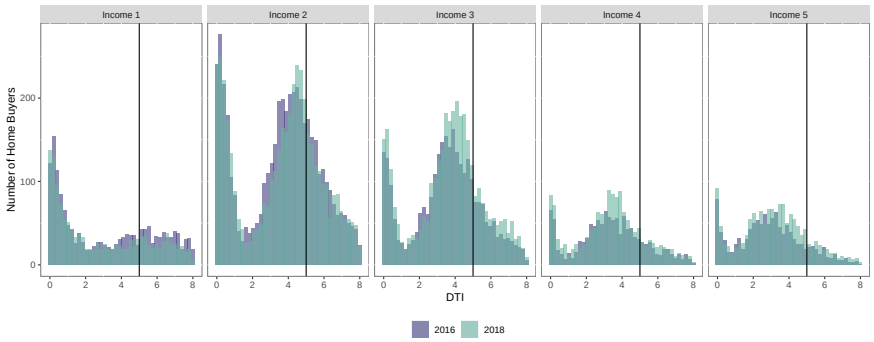
We report the relative choice sets available for two sets of households. The first is the median household, and the second is the top 10% households measured by their location in the income distribution. We define the relative choice as the affordable fraction of housing transactions within a month, 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 A1. 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 2 documents large heterogeneity in the effects of LTI limits on households' choice sets across income and housing products.



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 2 Effect of LTI Regulation on Housing Choice Sets

Last, we show evidence consistent with the LTI limit being binding for a subset of households. Figure 3 displays the distribution of Debt To Income (DTI) for households who purchased a property in 2016 (before the LTI limit, in blue) and in 2018 (after the LTI limit, in green) across five income groups (from lowest to highest income). We show how the LTI limit affects households' total debt because we do not observe directly their mortgage debt. Figure 3 shows that after the introduction of the LTI limit households in the middle of the income distribution (2^{nd} and 3^{rd} income groups), who purchased a property that year, exhibit bunching in their DTI, consistent with the LTI limit being a binding constraint.



Notes: These figures plot the distribution of Debt To Income (DTI) for households who purchased a property in 2016 (before the LTI limit, in blue) and in 2018 (after the LTI limit, in green) across five income groups (from lowest to highest income). The black vertical line in each figure identifies the LTI limit of 5 that was introduced in 2017. These distributions are bimodal because households can be first-time buyers, and hence more likely to rely on leverage to purchase a property, or switchers, and hence more likely to pay for most of the new property using the payment they receive from selling their old property.

Figure 3 Effect of LTI Regulation on Household Leverage

3 Model

The evidence in Table A3 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. The evidence in Figure 1 also raises 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 a 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 $i = 1, ..., \mathcal{N}$ potential buyers and sellers of residential real estate. The market participants are either households or investors and make a housing decision every year $t = 1, ..., T$.[15] Households own at most one residential 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. To reduce the dimensionality of the model, 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

[15] We use 2011-2018 for the estimation and 2010 for the initial housing allocation.

family size and age of the household head.[16]

3.1 Households

In every period t, a household i of type τ decides whether to rent or buy ($\mathcal{H} = \{0, 1\}$) and a housing product $j \in \mathcal{J}$. 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' housing options differ between renting and owning, and depend on whether they leave Oslo. Households leaving Oslo choose the outside option, so their housing choice is $d_{i,t} = 0$.[17] The housing choice of households who remain in Oslo is instead $d_{i,t} = \{j, h\}$. We assume that households can afford to spend up to 50 percent of their annual income on rent,[18] while their ownership options are constrained by LTI and LTV limits. Hence, we define the housing choice set for renting as $\mathcal{D}^0_{\tau,t} = \mathcal{J} \times \{0\}$, the housing choice set for ownership as $\mathcal{D}^1_{\tau,t} = \mathcal{J}_{\tau,t} \times \{1\}$, and the overall choice set as $\mathcal{D}_{\tau,t} = \{0\} \cup \mathcal{D}^0_{\tau,t} \cup \mathcal{D}^1_{\tau,t}$. The housing decision is defined as $d_{i,t} \in \mathcal{D}_{\tau,t}$, and moving occurs when a household changes its current housing choice, i.e., $d_{i,t} \neq d_{i,t-1}$.[19]

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

[16] To have a finite number of household types, we discretize the variables that define types. We divide income into five groups (all numbers are in NOK, with 1 USD corresponding to about 8 NOK in 2015): <400k, 400k-600k, 600k-800k, 800k-1,000k, >1,000k, 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 of household head into three groups: 18-34, 35-49, and above 50.

[17] We define the outside option as moving out of Oslo, reflecting households' mobility across cities. Conditional on moving, the average probability of leaving Oslo in our sample is 48.9%. While this probability is quite large, it is highly persistent over time within each household type. The mean variation within household types is only 3.8% of the the variation across household types, suggesting that the decision to leave Oslo is primarily driven by lifecycle mobility needs, such as education, employment, and migration, and that prices of properties in Oslo and outside Oslo follow similar trends. In the counterfactual analysis, we keep the utility of living outside Oslo fixed, assuming that factors affecting household utility, such as amenities, job opportunities, and the cost of living outside Oslo, remain unchanged. This allows us to quantify the costs and benefits of leverage regulation in Oslo while controlling for any spillover effects.

[18] This assumption is motivated by evidence in our data, as over 75% of households have an estimated user cost of housing below half of their income. We verified that assuming no affordability constraints related to rental costs has little impact on the results.

[19] For 7% of 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.

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:[20]

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

Note that tighter leverage limits can impact housing markets via two channels. First, by restricting households' choice set through the affordability constraints defined by equation (1). The importance of this channel, which we explicitly model, is supported by the evidence from Figure 2. Second, inducing lenders to adjust their mortgage interest rates due to the changes in housing demand. We do not model this second channel, not only because we lack households' mortgage-level data, but also because the evidence in Appendix C.1 suggests that this channel played almost no role when the LTI limit was introduced in Norway.

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:

$$\max_{\{d_{i,r} \in \mathcal{D}_{\tau,r}\}} \mathbb{E} \left[\sum_{r=t}^{T} \beta^{r-t} u \left(\mathbf{\Omega}_r, d_{i,r}, d_{i,r-1}, \varepsilon_{i,j,h,r} \right) | \mathbf{\Omega}_t, d_{i,t}, \varepsilon_{i,j,h,t} \right], \tag{2}$$

where β is the discount factor. Household utility is a function of the state variables $\mathbf{\Omega}_t$, which include: average price $P_{j,t}$, observed house and neighborhood characteristics (house size and neighborhood quality) $X_{j,t}$, unobserved house and neighborhood characteristics $\xi_{j,h,t}$, the LTI_t and LTV_t limits χ_t, household demographics (income, wealth, age, size) $Z_{i,t}$. It also depends on households' housing decisions in the current and last period $d_{i,t}$ and $d_{i,t-1}$, as well as latent demand $\varepsilon_{i,j,h,t}$ of household i for housing product j with the ownership status h. Since we allow for rich household-level heterogeneity and unobserved product-ownership-time level characteristics $\xi_{j,h,t}$, the remaining error $\varepsilon_{i,j,h,t}$ captures idiosyncratic shocks to housing choices and is assumed to be IID Type I Extreme Value distributed.

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

$$\begin{aligned} V(\mathbf{\Omega}_t, \varepsilon_{i,t}) = \max_{d_{i,t} \in \mathcal{D}_{\tau,t}} \{ & u \left(\mathbf{\Omega}_t, d_{i,t} \right) - \mathbb{1}_{\{d_{i,t} \neq d_{i,t-1}\}} F(Z_{i,t}, d_{i,t}, d_{i,t-1}) + \varepsilon_{i,t} \\ & + \beta \mathbb{E} \left[V(\mathbf{\Omega}_{t+1}, \varepsilon_{i,t+1}) | \mathbf{\Omega}_t, \varepsilon_{i,t}, d_{i,t} \right] \}, \end{aligned} \tag{3}$$

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

where $F(Z_{i,t}, d_{i,t}, d_{i,t-1})$ represents moving costs, and $\varepsilon_{i,t}$ represents the vector of choice-specific idiosyncratic shocks. 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 explicit affordability constraints or transaction costs. They trade freely each period. We index variables and parameters that apply to real estate investors by s. Given that, differently from the case of households, we cannot match each property to each individual investor, we sum up the holdings of all investors across each housing product j to one aggregate investor portfolio and model the rebalancing of that aggregate portfolio. As we do not observe the balance sheet of all investors individually, we impose no financial constraints in this rebalancing.

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 Koijen and Yogo (2019),[21] where investors set their optimal aggregate holdings as a function of housing product characteristics and a proxy for returns on housing.

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 $v^{\tau}_{d,t} = \overline{V}(Z_{i,t}, \chi_t, d_{i,t} = d)$ denote the choice-specific expected value function for a household with characteristics $Z_{i,t}$ and decision $d_{i,t}$. More explicitly, for households who stay in Oslo and choose $d_{i,t} = \{j, h\}$, the expected value function is $v^{\tau}_{j,h,t}$, while for those who leave Oslo the expected value function is $v^{\tau}_{0,t}$. To simplify the notation, we use τ to denote the household type $\tau(Z_{i,t})$. 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.

[21] The modeling choice can be justified with a mean-variance optimizing investor and i.i.d housing returns.

As standard in discrete choice models, a normalization is necessary to identify the vector of lifetime utilities $v^{\tau}_{d,t}$. In our case, we estimate a normalized lifetime utility $\tilde{v}^{\tau}_{d,t} = v^{\tau}_{d,t} - m^{\tau}_{t}$, where m^{τ}_{t} 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^{τ}_{t} in detail.

3.3.1 Households

We first charactere the households' decisions. Household i, when considering moving, will internalize the moving costs that transition their type to $\bar{\tau}$. The household chooses option d if $\tilde{v}^{\bar{\tau}}_{d,t} + \varepsilon_{i,d,t} > \tilde{v}^{\bar{\tau}}_{d',t} + \varepsilon_{i,d',t}\ \forall d' \neq d$. Conditional on moving to an inside option (i.e., for $d \neq 0$), the probability of a household of type $\bar{\tau}$ choosing housing product j with ownership h in period t is:

$$\Pr^{\bar{\tau}}_{j,h,t} = \frac{\exp(\tilde{v}^{\bar{\tau}}_{j,h,t})}{\sum_{d \in \mathcal{D}_{\bar{\tau},t} \setminus \{0\}} \exp(\tilde{v}^{\bar{\tau}}_{d,t})}. \tag{4}$$

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

$$L^{\text{prod}}_{i}(\tilde{v}) = \prod_{d \in \mathcal{D}_{\bar{\tau},t} \setminus \{0\}} \left(\Pr^{\bar{\tau}}_{d,t_{1,i}} \right)^{\mathbb{1}[d_{i,t_{1,i}} = d]}. \tag{5}$$

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

$$\Pr^{\bar{\tau}}_{0,t} = \frac{\exp(\tilde{v}^{\bar{\tau}}_{0,t})}{\sum_{d \in \mathcal{D}_{\bar{\tau},t}} \exp(\tilde{v}^{\bar{\tau}}_{d,t})}. \tag{6}$$

Let $t_{2,i}$ denote the period household i considers the outside option (conditional on moving). The likelihood contribution for this choice is:

$$L^{\text{out}}_{i}(\tilde{v}) = \Pr^{\bar{\tau}}_{0,t_{2,i}}{}^{\mathbb{1}[d_{i,t}=0]} (1 - \Pr^{\bar{\tau}}_{0,t_{2,i}})^{\mathbb{1}[d_{i,t}\neq 0]}. \tag{7}$$

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. A household who is currently in housing product j with ownership status h will choose to stay if:

$$v^{\tau}_{j,h,t} + \varepsilon_{i,j,h,t} > \max_{d \in \mathcal{D}_{\bar{\tau},t} \setminus \{j,h\}} [v^{\bar{\tau}}_{d,t} + \varepsilon_{i,d,t}] - \text{PMC}^{\bar{\tau}}_{i,t}, \tag{8}$$

where $\text{PMC}^{\bar{\tau}}_{i,t} = \bar{Z}'_{i,t}\gamma_{\text{pmc}}$ 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 $\bar{Z}_{i,t}$ and a vector of cost parameters γ_{pmc} to be estimated. By the definition of the normalized choice-specific value functions, we let $\tilde{v}^{\tau}_{d,t} = v^{\tau}_{d,t} - m^{\tau}_{t}$. Substituting this into the above equation gives:

$$\tilde{v}^{\tau}_{j,h,t} + \varepsilon_{i,j,h,t} > \max_{d \in \mathcal{D}_{\bar{\tau},t} \setminus \{j,h\}} [\tilde{v}^{\bar{\tau}}_{d,t} + \varepsilon_{i,d,t}] - (m^{\tau}_{t} - m^{\bar{\tau}}_{t}) - \text{PMC}^{\bar{\tau}}_{i,t}. \tag{9}$$

The term $(m^{\tau}_{t} - m^{\bar{\tau}}_{t})$ 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 $(m^{\tau}_{t} - m^{\bar{\tau}}_{t})$ 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, and the price of the property that the household is leaving $P_{d_{i,t-1},t}$.[22] Formally, we define it as:

$$m^{\tau}_{t} - m^{\bar{\tau}}_{t} = \text{FMC}^{\bar{\tau}}_{i,t} \gamma^{\bar{\tau}}_{i,\text{fmc}}. \tag{10}$$

We have two types of financial moving costs: in case of house purchase $\text{FMC}^{\bar{\tau}}_{i,t} = 0.05 \times P_{d_{i,t-1},t}$.[23] If instead a renter or a homeowner is moving to a rented property $\text{FMC}^{\bar{\tau}}_{i,t} = 0.02 \times P_{d_{i,t-1},t}$, capturing a combination of brokerage fees, potential loss of rental deposit, and other transaction fees. We let $\gamma^{\bar{\tau}}_{i,\text{fmc}} = \bar{Z}'_{i,t}\gamma_{\text{fmc}}$ to allow a marginal change in wealth to have a different impact on household utility depending on household characteristics. The probability that a household stays in its current property of type j with ownership status h at time t becomes:

$$\Pr^{\tau,\bar{\tau}}_{\text{stay},i,t} = \frac{\exp(\tilde{v}^{\tau}_{j,h,t})}{\exp(\tilde{v}^{\tau}_{j,h,t}) + \sum_{d \in \mathcal{D}_{\bar{\tau},t} \setminus \{j,h\}} \exp(\tilde{v}^{\bar{\tau}}_{d,t} - \text{FMC}^{\bar{\tau}}_{i,t}\gamma^{\bar{\tau}}_{i,\text{fmc}} - \bar{Z}'_{i,t}\gamma_{\text{pmc}})}. \tag{11}$$

The likelihood contribution of each household's sequence of move/stay decisions is:

[22] We follow Bayer et al. (2016) and express the financial moving costs as a function of the price of the property that the household is leaving (either selling it or leaving a rental), rather than the price of the new house the households is moving to. This allows us to estimate the model in two simple and tractable steps, rather than a single complex one.

[23] Housing 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%. Given that broker fees are typically between 1.5% and 3.5% of the transaction price, we set the total transaction cost to 5% of the selling house price.

$$L_i^{\text{stay}}(\tilde{v}, \gamma_{\text{fmc}}, \gamma_{\text{pmc}}) = \prod_{t=t_{0,i}}^{t_{0,i}+T_i} (P^{\tau,\bar{\tau}}_{\text{stay},i,t})^{\mathbb{1}[d_{i,t}=d_{i,t-1}]} (1 - P^{\tau,\bar{\tau}}_{\text{stay},i,t})^{\mathbb{1}[d_{i,t}\neq d_{i,t-1}]}. \tag{12}$$

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. While unrealistic, this assumption is motivated by the lack of data on investors' balance sheet, which prevents us from observing their leverage limits. We do expect however that on average investors' financial constraints will be captured in reduced-form by our estimates of their demand elasticity.

We assume investors' demand for product j at time t, expressed as the stock $\mathcal{S}^s_{j,t}$ of properties of type j at time t owned by investors, is an exponential function of rental yield $\frac{R_{j,t}}{P_{j,t}}$, defined as the ratio of rental price $R_{j,t}$ to transaction price $P_{j,t}$, property characteristics $X_{j,t}$, including districts and size fixed effects, and latent demand $\xi^s_{j,t}$:

$$\mathcal{S}^s_{j,t} = \exp\{\alpha^s \frac{R_{j,t}}{P_{j,t}} + \beta^s X_{j,t} + \xi^s_{j,t}\}, \tag{13}$$

where α^s is the elasticity of demand with respect to yield, β^s is the semi-elasticity of demand with respect to property attributes, and $\xi^s_{j,t}$ is unobserved demand for product j unrelated to investment returns or property attributes.

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 investors' housing portfolio.

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 supply part "# New Houses" represent the number of new properties that are built and put on the market for the first time at time t. This is the only exogenous component

of this market clearing condition, but given that we focus on a large and densely populated area such as Oslo, new constructions represent a very small fraction of the housing stock each year. In the demand part, switchers are homeowners who change residency, and entries are new households that exogenously enter the sample. These new entries represent for example young individuals who move out from their parents' home and enter the housing market. 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:

$$\underbrace{\mathcal{S}_{j,t}^{New} + (\mathcal{S}_{j,t-1}^{s} - \mathcal{S}_{j,t}^{s}) + \sum_{\tau} \sum_{i \in \tau} \left(1 - \Pr_{i,t}^{\tau,\bar{\tau}} \left(\{j,1\} \mid \{j,1\}\right)\right)}_{\text{Supply}} \tag{14}$$

$$= \underbrace{\sum_{\tau} \sum_{k} \sum_{i \in \tau} \Pr_{i,t}^{\tau,\bar{\tau}} \left(\{j,1\} \mid \{k,0\}\right) + \sum_{\tau} \sum_{k} \sum_{i \in \tau} \Pr_{i,t}^{\tau,\bar{\tau}} \left(\{j,1\} \mid \{k,1\}\right) + \sum_{\bar{\tau}} \mathcal{N}_{t}^{\bar{\tau}} \Pr_{j,h=1,t}^{\bar{\tau}}}_{\text{Demand}}.$$

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 last term in the supply is the number of properties sold by homeowners, where we sum up the probabilities that each homeowner belonging to type τ and owning product j does not remain in its property $1 - \Pr_{i,t}^{\tau,\bar{\tau}} \left(\{j,1\} \mid \{j,1\}\right)$, which gives the number of type τ homeowners that sell product j at time t. 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 (14). 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 (4); $\mathcal{N}_{t}^{\bar{\tau}}$ 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:

$$\Pr_{i,t}^{\tau,\bar{\tau}}(d_{i,t} = \{j,h\} \mid d_{i,t-1} = \{k,l\})$$
$$= \frac{\exp(\tilde{v}_{j,h,t}^{\bar{\tau}} - \text{FMC}_{i,t}^{\bar{\tau}} \widehat{\gamma}_{i,\text{fmc}}^{\bar{\tau}} - \bar{Z}_{i,t}' \widehat{\gamma}_{\text{pmc}})}{\exp(\tilde{v}_{k,l,t}^{\bar{\tau}}) + \sum_{d \in \mathcal{D}_{\tau,t} \setminus \{k,l\}} \exp(\tilde{v}_{d,t}^{\bar{\tau}} - \text{FMC}_{i,t}^{\bar{\tau}} \widehat{\gamma}_{i,\text{fmc}}^{\bar{\tau}} - \bar{Z}_{i,t}' \widehat{\gamma}_{\text{pmc}})}. \tag{15}$$

Note that the financial moving costs $\text{FMC}_{i,t}^{\bar{\tau}}$ are determined by $d_{i,t-1}$, as it is a function of the property price of the product that will be left. 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, and because properties up for rental can remain vacant. Investors own all the houses available for rent. We allow rental prices $R_{j,t}$ to adjust across counterfactuals using a simple hedonic rental pricing model described in Section 4.4.

In summary, house prices enter the market clearing condition via three channels. First, house prices determine financial moving costs as of equation (11), and accordingly higher prices reduce the likelihood that households move. Second, house prices determine investors' holdings of housing products based on equation (13), adjusting the denominator of investors' yield. Higher house prices will increase investors' cost of acquiring housing assets, hence reducing their yield and demand for those assets. The third channel is indirect and works via rental prices. As we will clarify in Section 4.4, rental prices also adjust in equilibrium as house prices and housing decisions change. More specifically, rental prices change as a function of the share of properties available for rent, relative to those owned and occupied by households. Intuitively, if higher house prices reduce households' homeownership rate relative to investors, this will lead to a larger supply of rental properties and hence a reduction in rental prices. This will affect the numerator of investors' yield as of equation (13). This will also affect households' indirect utility, as rental prices are used as a proxy for households' user cost of housing, as we describe in Section 4.3.

There is one last channel that we do not explicitly incorporate in the market clearing condition, which is households' expectations of future house prices. As we describe in Section 4.3, households form expectations over the evolution of house prices with an autoregressive process that depends on past prices, a common time trend, and housing product fixed effects. We find that, relative the common time trend and to housing product fixed effects, past prices have little explanatory power over the evolution of future prices. For this reason, and for computational simplicity, we do not allow future house price expectations to adjust via the market clearing condition.

4 Estimation

We estimate the model in four steps. First, we estimate $\tilde{v}$ from households' housing product choices. Second, we estimate the moving cost parameters $\gamma_{\text{fmc}}, \gamma_{\text{pmc}}$ 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 investors' portfolio problem and the hedonic rental price model.

4.1 First Stage: Value Functions

The closed-form solution that results from maximizing the likelihood of choosing a housing product or the outside option, conditional on moving,[24] with respect to $\tilde{v}$ is given by:

$$\widehat{\tilde{v}}^{\bar{\tau}}_{j,h,t} = \ln(\widehat{\Pr}^{\bar{\tau}}_{j,h,t}) - \frac{1}{\|\mathcal{D}_{\bar{\tau},t}\|} \sum_{d \in \mathcal{D}_{\bar{\tau},t}} \ln(\widehat{\Pr}^{\bar{\tau}}_{d,t}), \tag{16}$$

where $\widehat{\Pr}^{\bar{\tau}}_{j,h,t}$ is the empirical probability that a household of type $\bar{\tau}$ chooses housing product j with ownership status h at time t conditional on moving. $\|\mathcal{D}_{\bar{\tau},t}\|$ denotes the number of elements in the choice set of household type $\bar{\tau}$ at time t. 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 is that it mitigates potential small sample problems caused by having many household types.[25]

4.2 Second Stage: Moving Costs

Once we recover $\widehat{\tilde{v}}$, we find the parameters of the moving cost function $\gamma_{\text{fmc}}, \gamma_{\text{pmc}}$ that maximize the likelihood of the decision to stay or move, taking $\tilde{v}$ as given:

$$\max_{\gamma_{\text{fmc}}, \gamma_{\text{pmc}}} \sum_{i}^{N} \ln(L_i^{\text{stay}}(\widehat{\tilde{v}}, \gamma_{\text{fmc}}, \gamma_{\text{pmc}})). \tag{17}$$

The estimated $\widehat{\gamma}_{\text{fmc}}, \widehat{\gamma}_{\text{pmc}}$ can be used to recover the true choice-specific value functions $v^{\tau}_{j,h,t} = \tilde{v}^{\tau}_{j,h,t} + m^{\tau}_{t}$. Notice that $m^{\tau}_{t} - m^{\bar{\tau}}_{t}$ 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. We

[24] From Section 3.3.1 this likelihood function is $\sum_{i=1}^{N} \Big(\ln(L_i^{\text{prod}}(\tilde{v})) + \ln(L_i^{\text{out}}(\tilde{v})) \Big)$.

[25] We describe the details of the kernel smoothing in Appendix D.

hence set $m_t^\tau = A^\tau \gamma_{\text{fmc}}^\tau$, where A^τ is the median wealth of a type τ household. This recovers the wealth effect on household utility $v_{j,h,t}^\tau$ for all household types.

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 the following autoregressive processes:

$$v_{j,h,t}^\tau = \psi_{0,j,h}^\tau + \sum_{l=1}^{2} \psi_{1,l} v_{j,h,t-l}^\tau + \sum_{l=1}^{2} \psi_{2,l} P_{j,t-l} + \psi_{3,j,h}^\tau t + \omega_{j,h,t}^\tau, \tag{18}$$

$$P_{j,t} = \phi_{0,j} + \sum_{l=1}^{2} \phi_{1,l} P_{j,t-l} + \phi_{2,j} t + \rho_{j,t}. \tag{19}$$

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:

$$u_{j,h,t}^\tau = v_{j,h,t}^\tau - \beta \mathbb{E}\left[\ln\left(e^{v_{j,h,t+1}^\tau} + \sum_{d \in \mathcal{D}_{\bar{\tau},t+1} \setminus \{j,h\}} e^{v_{d,t+1}^{\bar{\tau}} - \text{FMC}_{i,t+1}^{\bar{\tau}} \widehat{\gamma}_{i,\text{fmc}}^{\bar{\tau}} - \bar{Z}_{i,t+1}' \widehat{\gamma}_{\text{pmc}}} \right) \mid s_{i,t}, d_{i,t} = \{j,h\} \right], \tag{20}$$

where, in practice, $s_{i,t}$ includes all the variables on the right-hand side of equations (18) and (19), 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 (20). 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. We use the draws of house prices to determine housing wealth, which gives households' type in the next period τ_{t+1}. For each draw, we calculate a per-period flow utility $u_{j,h,t}^\tau$ using equation (20). The average across draws is the simulated $u_{j,h,t}^\tau$.

We follow Bayer et al. (2016) to estimate the determinants of flow utility of household type τ for housing product j with ownership status h at time t:

$$u_{j,h,t}^\tau = \alpha_\tau^u + \alpha_t^u + \alpha_h^u + X_{j,t} \alpha_x^u + \alpha_r^u R_{j,t} + \xi_{j,h,t}^\tau, \tag{21}$$

where α_τ^u are household-type fixed effects, α_t^u are year fixed effects, α_h^u represents the utility of homeownership, $X_{j,t}$ includes district quality and house size dummies, and $\xi_{j,h,t}^\tau$

captures type-specific unobserved house and district attributes. We also include in households' flow utilities the rental price $R_{j,t}$ to proxy for the user cost of housing.[26] To address the endogeneity of the user cost of housing, we exploit the estimated $\widehat{\gamma}^{\tau}_{\text{fmc}}$ 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 (21) as:

$$u^{\tau}_{j,h,t} + \widehat{\gamma}^{\tau}_{\text{fmc}} R_{j,t} = \alpha^{u}_{\tau} + \alpha^{u}_{t} + \alpha^{u}_{h} + X_{j,t}\alpha^{u}_{x} + \xi^{\tau}_{j,h,t}. \tag{22}$$

We estimate the parameters $\alpha^{u}_{\tau}, \alpha^{u}_{h}, \alpha^{u}_{t}, \alpha^{u}_{x}$ with a linear regression.

4.4 Fourth Stage: Investors' Portfolio Choice

We estimate α^{s}, β^{s}, and recover latent demand $\xi^{s}_{j,t}$, in the investors' model such that: (i) the investors' demand function defined by equation (13) holds, (ii) households' decisions $\Pr^{\tau,\bar{\tau}}_{i,t}(d_{i,t}|d_{i,t-1})$ defined by equation (15) hold, and (iii) the market clearing conditions defined by equation (14) are satisfied. To achieve this, we first use the households' decision model to obtain model-implied housing demand and supply by all types of households $\widehat{\Pr^{\tau,\bar{\tau}}_{i,t}}(d_{i,t}|d_{i,t-1})$. Next, we use the market clearing conditions to back out the model-implied housing stock held by investors $\widehat{\mathcal{S}}^{s}_{j,t}$, using the observed newly constructed houses, the number of new households entered into Oslo, and the stock holding of investors from the previous period. Finally, taking the logarithm of equation (13), we estimate the parameters of the investors' demand function as follows:

$$\ln(\widehat{\mathcal{S}}^{s}_{j,t}) = \alpha^{s}\frac{R_{j,t}}{P_{j,t}} + \beta^{s}X_{j,t} + \xi^{s}_{j,t}. \tag{23}$$

To allow for changes in rental prices across counterfactuals, given that we do not impose a market clearing condition for the rental market, we estimate the following hedonic rental pricing model:

$$R_{j,t} = \alpha^{r}O_{j,t} + \alpha^{r}_{j} + \alpha^{r}_{t} + \varepsilon^{r}_{j,t}, \tag{24}$$

where rental prices $R_{j,t}$ depend on the fraction $O_{j,t}$ of the housing product j at time t that is occupied by homeowners, reflecting the availability of housing product j in rental markets, product fixed effects α^{r}_{j}, time fixed effects α^{r}_{t}, and an error term $\varepsilon^{r}_{j,t}$. We keep the estimated parameters fixed across counterfactuals, but allow $O_{j,t}$ to vary.

[26] We 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.

4.5 Identification

We exploit primarily three sources of variation in the data that provide identification of the parameters of households' indirect utility and moving cost functions, as well as of investors' demand functions. First, every period a new set of houses enters the real estate market in Oslo, ranging from 2,500 to 5,500 units across all districts. These new properties are the result of urban planning decisions taken several years before their entry into the market, mitigating the concern of possible correlation with unobserved time-location-specific shocks. Second, every period the set of households making housing decisions changes, as new households are formed, some old households exit the sample, and households change their type due to changes in income, wealth, age, and composition. More specifically, on average every year 6.9% of households are newly formed, and 38.5% of households change type due to changes in income (23.2%), wealth (13.8%), age (5.1%), and composition (6.3%). We consider these household changes as orthogonal relative to unobservables in our model. Last, within our sample period in 2017 the LTI limit of 5 was introduced, providing substantial variation in households' choice set, with heterogeneous effects across households and housing products, as documented in Section 2.4.

5 Results

5.1 Moving Costs

We find that the average moving cost is 126,700k NOK (∼10,136 USD), while the average income is 541,400k NOK (∼43,312 USD). As a result, the median of the ratio of moving cost to income is 22%. To estimate the financial and psychological moving cost parameters γ_{fmc} and γ_{pmc}, we use the likelihood function (12), which is based on households' mobility choices. Table A4 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 costs depend on households' income, age group, and family size.[27] The takeaways from Table A4 are as follows. First, and as expected, financial moving costs matter less for high-income households. Second, psychological moving costs also decrease with household income. Third, older and larger households face high psychological moving costs.

[27] The characteristics are defined in footnote 16.

5.2 Flow Utilities and Willingness To Pay

We now regress the flow utility $u^{\tau}_{j,h,t}$ on property and neighborhood characteristics, excluding user costs in line with equation (22). The coefficients presented in Table A5 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 our preferred specification, 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 is preferred to renting at an increasing rate with income.

Panel B of Table 2 reports households' willingness to pay for housing attributes by income. We scale the estimates by dividing the estimated coefficients from equation (22) by the marginal utility of wealth $\hat{\gamma}^{\tau}_{\text{fmc}}$. 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. For example, an average household with an annual income between 600,000 and 800,000 NOK (∼48,000-64,000 USD) is willing to pay annually 53,200 NOK (∼4,200 USD) to live in a high-quality neighborhood relative to a middle-quality. The willingness to pay for homeownership is increasing across income groups. Overall, households have higher willingness to pay for smaller properties, again increasing with income, but as shown in Table A6 this is mostly driven by renters. For homeowners, we find that households in the two lowest income groups are willing to pay increasingly more for larger properties. Those in the third and fourth income groups 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.[28]

We derive the willingness to pay taking into account the affordability constraints that bound households' choice sets. However, due to lack of data on consumption and saving, we

[28]These results are comparable to the willingness-to-pay 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 neighborhood amenities.

Table 2 Willingness to Pay for Housing Attributes

	Income Groups				
	1	**2**	**3**	**4**	**5**
A. Household Characteristics					
Home-Owners	43.7%	64.3%	74.4%	78.9%	82.3%
Less Than 35 Years Old	36.5%	34.9%	26.1%	17.6%	10.1%
B. Willingness to Pay					
High Quality	31.3	41.1	53.2	70.6	120.2
Low Quality	11.5	4.4	−4.3	−17.0	−53.0
Three Rooms	−16.6	−25.7	−36.9	−53.1	−99.2
Four Rooms and Above	−40.6	−65.0	−95.1	−138.6	−262.4
Ownership	102.7	126.0	154.7	196.2	314.4
C. Willingness to Pay to Income					
High Quality	10.8	8.4	7.8	8.0	9.8
Low Quality	4.0	0.9	−0.6	−1.9	−4.3
Three Rooms	−5.7	−5.3	−5.4	−6.0	−8.1
Four Rooms and Above	−14.1	−13.3	−14.0	−15.7	−21.5
Ownership	35.6	25.8	22.8	22.2	25.7

Notes: This table reports household characteristics and the average willingness to pay for housing attributes of households with different income levels. Income groups are defined based on yearly household income in NOK of <400k (group 1), 400-600k (group 2), 600-800k (group 3), 800-1,000k (group 4), >1,000k (group 5). Panel A reports the percentage of households that are homeowners (Home-Owners), and whose household head is less than 35 years old. In Panel B, the first (second) row shows how much more an average household is willing to pay in thousands of NOK for being in a High-Quality (Low-Quality) district relative to a Middle-Quality district every year. The third (fourth) row shows how much an average household in each 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.

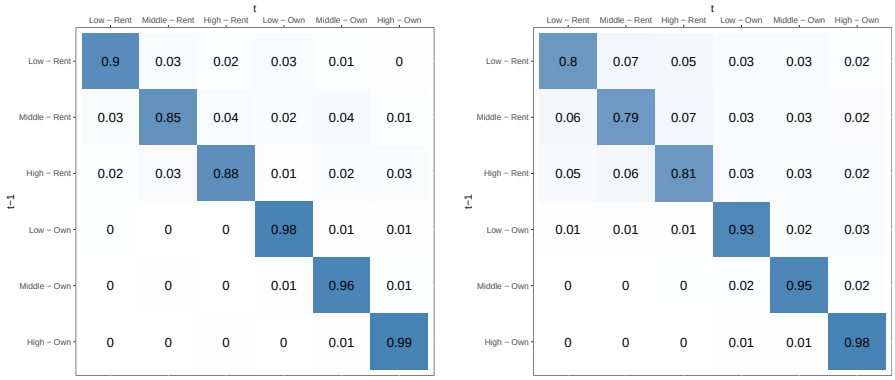
are not able to explicitly incorporate a budget constraint that bounds households' ability to spend every period on housing consumption. Therefore, the willingness to pay is based on the assumption that housing consumption is unconstrained by the per-period budget limit. This assumption is empirically supported by the result that the estimated user cost of housing is well below households' income.

5.3 Investors' Portfolio Choice

We estimate equations (23), (24) by ordinary least squares, and report the results in Table A7. We find that one standard deviation increase in rental yield increases the stock of housing held by investors by 34.9%, and that one standard deviation increase in the home-ownership rate for a given housing product increases its rental price by 6.7%.

5.4 Model Fit

We evaluate the fit of the model comparing the observed and model-predicted household mobility patterns. As reported in Figure 4, the model-predicted moving probabilities match well with the sample observations, although the predicted moving probabilities for renters are higher than the observed values.



Notes: These figures plot probabilities of moving, averaged across time within our sample period. The left vertical axis refers to locations and ownership of households in year $t-1$, while the top horizontal axis refers to locations and ownership of household in the following year t. Locations can be of three quality levels ("Low", "Middle", "High"), and within these locations households can be renters ("Rent") or homeowners ("Own"). For example, the top-left cell of each figure is the probability that a household who rents in a low-quality district in year $t-1$ remains in a low-quality district as renter in year t. The left figure represents the observed moving probabilities, while the right figure represents the model predicted probabilities. For ease of interpretation, these are all probabilities conditional on not moving out of Oslo.

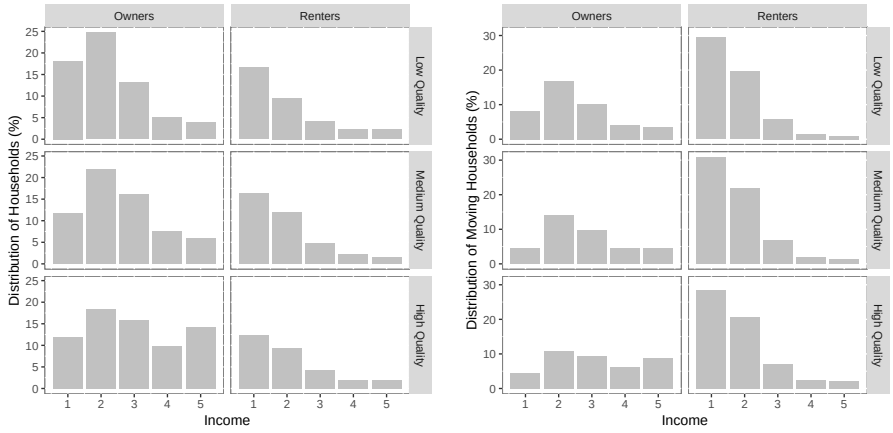
Figure 4 Observed vs Model-Implied Mobility Patterns

6 Counterfactuals

We simulate counterfactual scenarios to quantify the distributional effects of leverage limits, and to investigate whether three commonly used housing policies can mitigate the regressive effects of these limits. First, we simulate a scenario for 2017 in which the LTI is three instead of five. We focus on the counterfactual house prices, household debt, and mobility patterns across the income distribution. Second, we simulate a progressive subsidy aimed at offsetting the regressive impact of tighter LTI limits. We implement it reducing the distance in willingness to pay for house and district attributes and homeownership of the first four income groups relative to the top income group by 50%. This exercises is a way to proxy for housing policies, such as schooling vouchers and public housing assistance programs, that encourage low-income households to become homeowners and live in high-quality neighborhoods. Third, we simulate an expansion of housing supply, and the introduction of regulated

rental prices, documenting how these impact homeownership rates and housing market segmentation across income groups.

Before showing the counterfactual results, we present in Figure 5 the 2017 baseline distribution, by homeownership status, of households across house quality and income. The 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.



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 income and homeownership status.

Figure 5 Stocks and Flows of Households across Quality and Income

There are two takeaways from the left figure. First, across all quality levels, the proportion of owners relative to renters increases monotonically with income. This suggests low-income households are primarily young people who, for instance, have not yet accumulated enough wealth and income to afford a house. This means that real estate investors' portfolios are tilted towards properties in low-quality districts. Second, the largest proportion of high-income households is in high-quality districts, and low-income households dominate low-quality districts, consistent with housing market segmentation across the income distribution.

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 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 3 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% (73.6%) 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 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 68.9% of the housing market before the introduction of the LTI, which then dropped by 17.4%.

Moreover, in the baseline 14.2% (23.3%) 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 25.6% drop in their access to high-quality housing as a result of a tighter LTI.

The other two rows of the top panel of Table 3 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 renters are more likely to move than low-income ones. Third, a stricter LTI limit has small impact on households' probability of moving, because the baseline probability of moving is low.

The second panel of Table 3 focuses on households that move. The probability of renters becoming homeowners, mostly representing first-time buyers, is 3.8% (36.0%) for the lowest (highest) income group. A tighter LTI limit reduces it by 37.4% (0.0%). The probability

Table 3 Effect of LTI Change on Choice Sets and Moving Probabilities

		Income Groups				
	LTI	**1**	**2**	**3**	**4**	**5**
All households						
Owners' Share of Total Products in Choice Set	Base	52.9	62.2	68.9	73.3	73.6
	Δ	−14.3%	−18.6%	−17.4%	−15.2%	−1.0%
Owners' Share of High Quality Products in Choice Set	Base	14.2	16.9	19.1	21.3	23.3
	Δ	−12.5%	−18.4%	−25.6%	−20.8%	−3.2%
Owners' Moving Probability	Base	6.7	6.8	5.8	5.6	5.4
	Δ	1.6%	1.6%	1.0%	0.6%	0.9%
Renters' Moving Probability	Base	15.6	35.6	35.5	33.2	30.7
	Δ	0.0%	−1.4%	−0.7%	−0.4%	0.2%
Movers						
From Renting to Owning	Base	3.8	13.5	23.9	31.3	36.0
	Δ	−37.4%	−43.4%	−11.3%	−4.7%	0.0%
From Low to High Quality	Base	5.1	6.5	7.6	9.5	10.9
	Δ	−1.5%	−2.0%	−6.1%	−4.7%	−0.5%

Notes: This table reports the mobility of households for each income group 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.

of climbing the housing quality ladder is almost 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 up to 6.1% for middle-income households.[29]

In Table 4, we show how stricter LTI limits reduce household debt. Motivated by the evidence in Table 3, our focus is on moving households.[30] Specifically, we calculate the lower and upper bounds of total debt across five income groups, defined as the sum of total debt

[29]In Table A8 we report the same table as Table 3 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.

[30]The corresponding results for all households are in Table A9 in the Appendix.

for all moving households within the income group. For the lower bound, we assume that households first use all of their wealth to cover the house investment and the moving costs and then borrow whatever is left. For the upper bound, we assume that households finance the new property with the housing wealth from the sale of their old property (if they have one) and finance everything else with debt.

We observe the largest debt reduction for the second income group, with a drop between 27.7% and 33.9%. For the lowest and the highest income groups, stricter LTI limits reduce total debt between 14.6% and 18.3% and between 2.5% and 2.9%, respectively. The finding that households in the middle of the income distribution are the most affected by a reduction in LTI is in line with the extensive choice set contraction reported in Table 3.

Table 4 Effect of LTI Change on Aggregate Household Leverage

		Income Groups				
	LTI	**1**	**2**	**3**	**4**	**5**
Movers						
Lower Bound Debt	Base	1.9	4.9	3.7	1.8	2.0
	Δ	−18.3%	−33.9%	−15.2%	−10.6%	−2.9%
Upper Bound Debt	Base	2.9	7.3	5.6	2.8	3.1
	Δ	−14.6%	−27.7%	−12.2%	−8.2%	−2.5%

Notes: This table presents the sum of total housing debt (in billions of NOK) of moving households across five income groups for the baseline scenario (Base) and the percentage changes under the counterfactual scenario of changing LTI limits (Δ). The top two rows reports the lower bound assumption of debt, where households use all their wealth to finance their home purchase and financial moving costs. The bottom two rows reports the upper bound assumption of debt, where households use only housing wealth to finance their home purchases, and finance the rest by debt.

In Table 5, we report changes in equilibrium prices between the baseline and the counterfactual LTI. A more stringent LTI reduces prices by up to 5.1%. Small sized and low-quality properties, the primary target of low- and middle-income households, drop the most in value. Two channels explain changes in equilibrium prices under counterfactual LTI. As affordability constraints tighten, marginally affordable houses drop out of households' choice sets, shifting their demand toward less expensive properties (intensive margin) or causing them to exit the ownership market entirely (extensive margin). These two effects have counteracting impacts on the prices of less expensive properties. As the results suggest, the extensive margin dominates, leading to a relatively larger negative price impact of stricter LTI limits on smaller properties and homes in lower-quality districts. This result has two distributional implications. On the one hand, it is progressive as it makes the less expensive houses relatively more affordable, favoring more financially unconstraint low-income first-time buyers.

On the other hand, it is regressive as homeowners of less expensive properties, who are typically low-income households, see their housing wealth depreciate relatively more.

Table 5 Effect of LTI Change on Prices

		LTI = 5	Δ LTI
House Size	Small	3.2	−5.1%
	Medium	4.2	−2.0%
	Large	6.3	−3.4%
District Quality	Low	3.1	−4.8%
	Medium	3.4	−4.6%
	High	4.3	−4.4%

Notes: This table reports the average house prices (in millions NOK) in 2017. The second column reports prices under the baseline LTI of 5, while the third column reports the percentage change in prices when the LTI is reduced to 3. The first panel reports the average house prices for small (1-2 bedrooms), medium (3 bedrooms), and large (4+ bedrooms), while the second panel is for low, medium, and high-quality districts.

6.2 Housing Subsidy

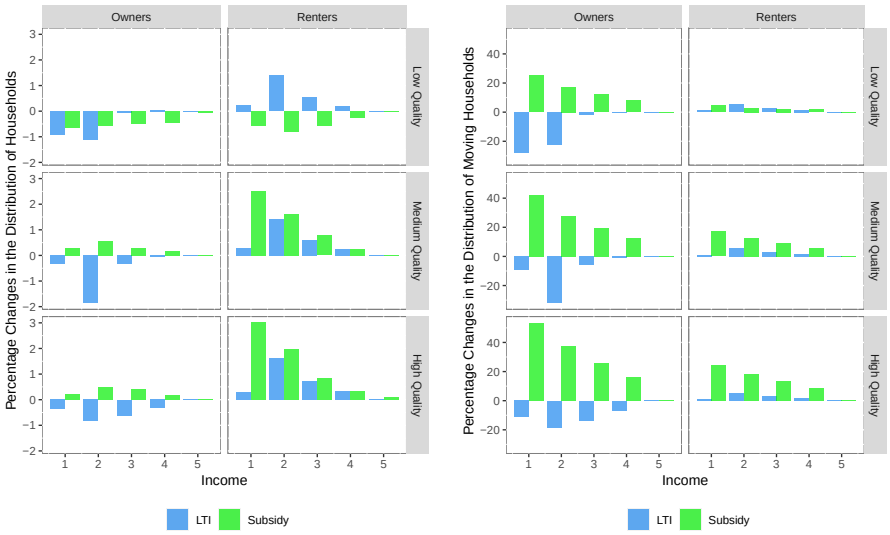
The second counterfactual quantifies the progressive effects of a housing subsidy, as well as the importance of preferences in residential choices. Progressive housing subsidy programs, aimed at improving housing affordability for low-income households in high-quality areas, have been implemented in several countries in various formats.[31] We simulate a housing policy that encourages low-income households to become homeowners or move to high-quality neighborhoods, and is equivalent to a yearly transfer of 47,000 NOK (3,800 USD) for the lowest income group, declining progressively as income increases. In practice, we simulate a scenario where we proportionally subsidize households' willingness-to-pay for ownership and housing attributes, reducing the distance between the willingness-to-pay of the bottom four income groups and the top one by 50%.[32] We do not allow however this subsidy to impact the set of housing options that households can afford. By doing so, we are able to measure the extent to which segmentation in housing markets is explained by preference

[31] Notable examples include mobility programs that help low-income households move to better districts, such as the Housing Choice Voucher and Moving to Opportunity programs in the United States; progressive rental subsidies for low-income households with higher subsidies in more expensive areas, as seen in Germany, France, and Finland; or subsidies that cover high-quality housing if certain criteria are met, as in the Netherlands, Sweden, and Canada. There are also homeownership programs that progressively benefit lower-income groups, such as the First-Time Homebuyer Program in the United States, the First Home Savings Account in Canada, and housing grants in Singapore.

[32] We obtain the willingness-to-pay 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.

heterogeneity, holding fixed their affordability constraint. In Figure 6 we compare the effects of this counterfactual, labeled as "Subsidy", to the effects of the first counterfactual, labeled as "LTI". 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 5.

In line with Table 3, both figures show how tighter LTI limits reduce 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. We find two main effects when we provide a housing subsidy. First, as the right panel of Figure 6 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 6 shows, households are reallocated from low- to high-quality neighborhoods, both for owners and renters. These effects materialize because the change in willingness to pay increases the benefit of being homeowners and living in high-quality districts.



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 two different counterfactuals. The blue bars represent the case of tighter LTI limits, and the green bars represent the case of a progressive housing subsidy.

Figure 6 Changes in Stocks and Flows across Quality and Income: Housing Subsidy

The overall takeaway of this comparison is the following. On the one hand, more stringent LTI limits can have regressive effects and increase segmentation in housing markets, as they

reduce homeownership rates primarily for low-income households. However, this does not prevent low-income households from living in high quality districts through the rental market. This is consistent with the results in Table 3, which show that more stringent LTI limits have a large impact on the transition from renting to owning, but only a moderate impact on the transition from low- to high-quality districts. On the other hand, providing housing subsidies can have a countervailing effect, as it encourages household mobility, especially for low-income households, and creates an incentive to move to high-quality neighborhoods, both for owners and renters. Taken together, these findings quantify the importance of preference heterogeneity, and show that progressive housing subsidies can effectively mitigate the unintended segmentation effects of LTI regulation.

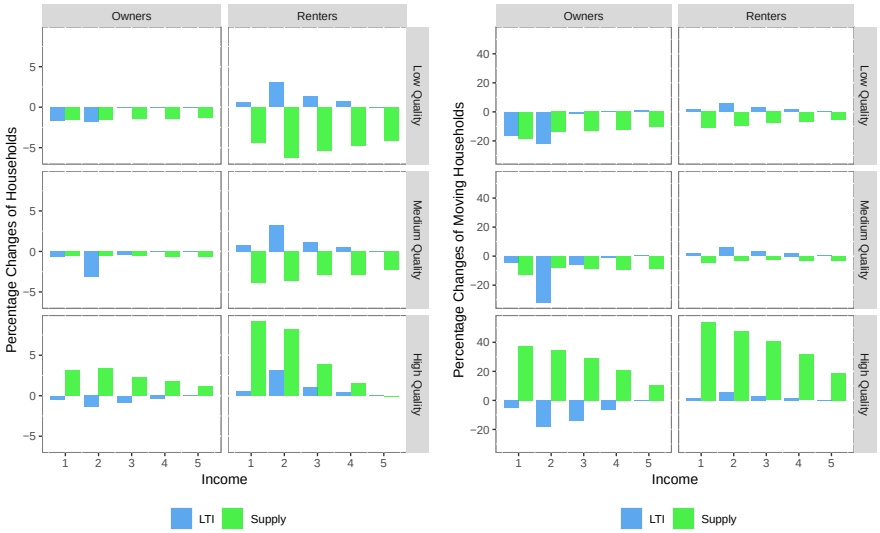
6.3 Housing Supply

As an alternative way to facilitate households' access to high-quality housing and homeownership, we simulate a scenario where we increase the housing stock in high-quality districts. We introduce 5,000 new housing units, allocated proportionally to the existing stock of small, medium, and large properties in high-quality neighborhoods. This is a relatively large increase in housing supply, as during our sample period the number of new housing units built annually, across all sizes and districts, ranges between 2,500 and 5,500. We constrain these new housing units to be only available to households rather than investors, to focus on the distributional effects across households. The resulting changes in the distribution of all households and moving households are shown in Figure 7, with the LTI counterfactual included for comparison. We observe a shift of households moving from low-quality and middle-quality to high-quality districts, with a stronger effect among lower-income households. This implies that an increase in housing supply in high-quality districts mitigates housing segmentation across income groups.

6.4 Rent Regulation

Finally, we consider a rent control policy that imposes caps on rental prices. Rent controls have been widely adopted across countries to protect tenants' rights and ensure housing affordability, particularly targeting low-cost properties and low-income households.[33] In

[33]For example, the Netherlands has a point-based rent control system that sets rental price caps in the regulated sector, mostly affecting small and low-quality properties. Austria applies rent regulation to specific segments, such as older buildings and those with low construction quality, with strict rent caps. Denmark uses cost-based rent control for older properties, where rents in regulated cities are calculated based on landlords' maintenance and ownership costs. Several French cities enforce rental caps, where rents cannot exceed more than 20% of the reference rent determined by property type, size, and location.



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 two different counterfactuals. The blue bars represent the case of tighter LTI limits, and the green bars represent the case of increasing 5,000 new housing units in high quality districts.

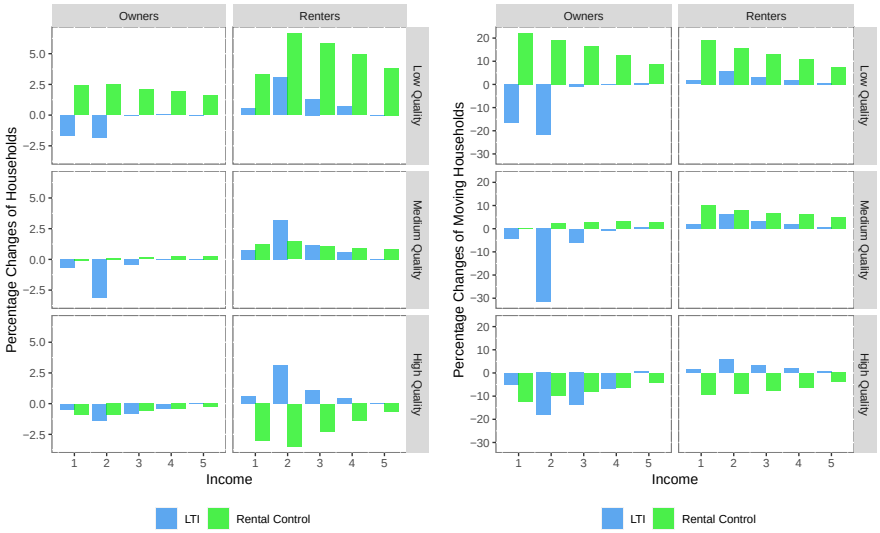
Figure 7 Changes in Stocks and Flows across Quality and Income: Housing Supply

Oslo, the rent-to-price ratio in 2017 ranged from 2.8% to 6.7% across different property types, with higher ratios observed in low-quality districts. We implement a rent-to-price ratio cap of 4%, which primarily affects properties in low-quality and some middle-quality districts, leading to lower rental prices in those areas under this counterfactual policy.

The results are presented in Figure 8. Two key takeaways emerge. First, renters move from high-quality to low-quality districts, because the reduction in rental prices for the latter areas is relatively larger. This is primarily driven by low-income households, as they have lower willingness to pay for district quality. Second, the rental cap reduces investors' yield relatively more in low-quality neighborhoods, which in turn reduces investors' demand for those properties. This implies that more households, especially the low-income ones, will be able to afford purchasing those houses. As a result, while rent control targeting low-quality properties improves housing affordability for lower-income groups, it also reinforces segmentation in housing markets across income groups.

7 Conclusion

We develop and estimate a structural model of housing demand and supply to quantify the distributional effects of leverage regulation, and to investigate whether three commonly used



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 two different counterfactuals. The blue bars represent the case of tighter LTI limits, and the green bars represent the case of setting rent-to-price ratio at 4%.

Figure 8 Changes in Stocks and Flows across Quality and Income: Rent Regulation

housing policies can mitigate the regressive effects of these limits. 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 counterfactual exercises. We first impose a tighter leverage limit relative to the baseline, which delivers the following results. First, we document that tighter limits are effective at reducing household debt and house prices. Second, 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 25.6% 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 up to 43.5% for low-income households but does not affect the richest. Overall, these findings highlight the inequality aspect of mortgage regulation.

To offset these regressive effects, we simulate a scenario where households receive a yearly housing subsidy equivalent to up to 47,000 NOK (3,800 USD) for the lowest income group, declining progressively as income increases. We implement this simulation by shifting households' willingness to pay for property and district characteristics 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. We also simulate the effects of an expansion of housing supply in high-quality districts, and of rent controls that are mostly binding in low-quality neighborhoods.

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. forthcoming at *Econometrica*.

Attanasio, O. P., R. Bottazzi, H. W. Low, L. Nesheim, and M. Wakefield (2012). Modelling the demand for housing over the life cycle. *Review of Economic Dynamics 15*(1), 1–18.

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. (2021). Leverage regulation and market structure: An empirical model of the U.K. mortgage market. *Journal of Finance LXXVI*(6), 2997–3053.

Benetton, M., A. Gavazza, and P. Surico (2025). Mortgage pricing and monetary policy. *American Economic Review 115*(3), 823–863.

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.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2024). Beyond the balance sheet model of banking: Implications for bank regulation and monetary policy. *Journal of Political Economy 132*(2), 616–693.

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.

Cocco, J. (2005). Portfolio choice in the presence of housing. *Review of Financial Studies 18*(2), 535–567.

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.

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

Fagereng, A., M. B. Holm, B. Moll, and G. Natvik (2021). Saving behavior across the wealth distribution: The importance of capital gains. Working Paper.

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.

Favilukis, J., P. Mabille, and S. Van Nieuwerburgh (2023). Affordable housing and city welfare. *Review of Economic Studies 90*(1), 293–330.

Fonseca, J., L. Liu, and P. Mabille (2024). Unlocking mortgage lock-in: Evidence from a spatial housing ladder model. Working Paper.

Geanakoplos, J. (2009). The leverage cycle. *NBER Macroeconomics Annual 2009 24*, 1–65.

Gomes, F. and A. Michaelides (2005). Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance LX*(2), 869–905.

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.

Kaplan, G., K. Mitman, and G. L. Violante (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy 128*(9), 3285–3345.

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.

Mabille, P. (2023). The missing homebuyers: Regional heterogeneity and credit contractions. *Review of Financial Studies 36*(7), 2756–2796.

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.

Ouazad, A. and R. Rancière (2018). City equilibrium with borrowing constraints: Structural estimation and general equilibrium effects. *International Economic Review 60*(2), 721–749.

Peng, Y. (2023). 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.

Piazzesi, M. and M. Schneider (2016). Housing and macroeconomics. *Handbook of Macroeconomics 2*, 1547–1640.

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.

Solheim, H. and B. H. Vatne (2013). Measures of household credit risk. Economic Commentaries Financial Stability No. 8 - 2013, Norges Bank.

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.

Tzur-Ilan, N. (2023). Adjusting to macroprudential policies: Loan-to-value limits and housing choice. *Review of Financial Studies 36*(10), 399–4044.

van Bekkum, 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*.

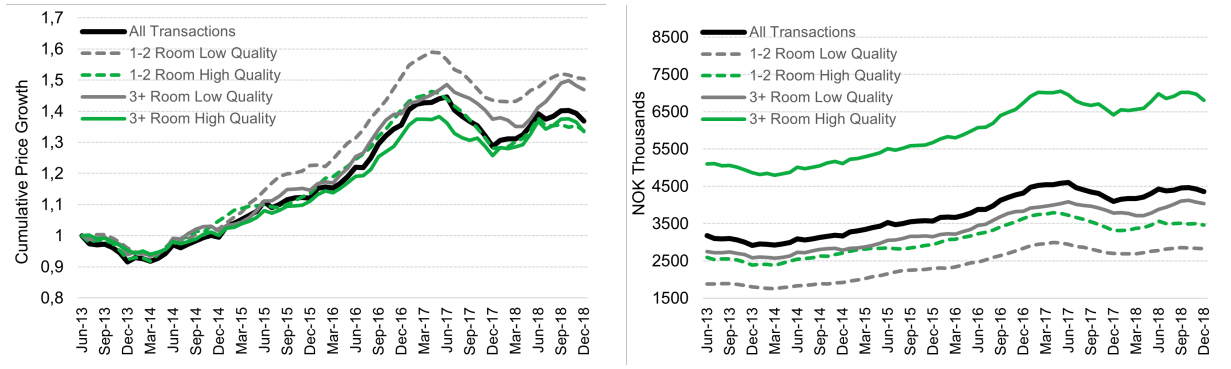
Online Appendix

Appendix A Supplemental Tables and Figures

Table A1 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 first column reports the total variance, the second column the variance between housing products, the third column the variance within housing products, while the last two columns report the percentage of the total variance explained by between and within variation. The 45 housing products are constructed as follows: we first define the collection of 15 districts as the set $\mathcal{D} = \{1, 2, ..., 15\}$. 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 apartments with one or two rooms, 3 includes apartment with three rooms, and 4^{+} 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 $\mathcal{J} = \{1, 2, ..., 45\}$ housing products.



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 A1 Effect of LTI Regulation on House Prices over Time

Table A2 House Price Regression

	Dependent variable: Log Price				
	(1)	(2)	(3)	(4)	(5)
House size (m^2)					0.01***
					(0.0000)
2011	0.09***	0.09***	0.09***	0.10***	0.10***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
2012	0.20***	0.20***	0.20***	0.20***	0.20***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
2013	0.23***	0.23***	0.24***	0.24***	0.25***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
2014	0.26***	0.26***	0.27***	0.26***	0.28***
	(0.004)	(0.003)	(0.004)	(0.003)	(0.002)
2015	0.39***	0.39***	0.39***	0.39***	0.40***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
2016	0.54***	0.54***	0.55***	0.55***	0.57***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
2017	0.60***	0.60***	0.62***	0.61***	0.64***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
2018	0.63***	0.62***	0.64***	0.63***	0.65***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
2019	0.67***	0.66***	0.68***	0.66***	0.68***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
2020	0.70***	0.71***	0.72***	0.71***	0.73***
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)
Month FE (12)	Yes	Yes	Yes	Yes	Yes
Room FE (3)	No	Yes	No	No	No
District FE (15)	No	No	Yes	No	No
Room-District FE (45)	No	No	No	Yes	No
Zip-code FE (406)	No	No	No	No	Yes
Nobs	312,020	312,020	312,020	312,020	312,020
R^2	0.22	0.47	0.38	0.66	0.80
Adj. R^2	0.22	0.47	0.38	0.66	0.80

Notes: This table reports the results from regressing the natural logarithm of house prices on time-series and product fixed effects, as well as house size measured in square meters (m^2). The sample covers the period from 2010 (the reference year) to 2020. The fixed effects include month fixed effects (12), number of rooms (3), district (15), number of rooms by district (45), and zip code (406 in total). The significance levels are given by *p<0.1; **p<0.05; ***p<0.01.

Table A3 Housing Quality Regressions

	Price	Income	Wealth	Age	Children
Baseline	5,191.56***	420.06***	110.60***	39.43***	0.74***
	(275.76)	(14.18)	(14.58)	(1.56)	(0.15)
Low Quality	−1,670.54***	−42.20**	−39.66*	4.26*	0.53**
	(356.00)	(18.30)	(18.83)	(2.02)	(0.20)
High Quality	1,625.81***	67.62***	72.09***	5.97**	0.45**
	(369.97)	(19.02)	(19.57)	(2.10)	(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 financial wealth, age and number of children.

Table A4 Financial and Psychological Moving Cost

Financial Moving Cost	
Constant	0.006***
	(0.00)
Income	-0.004***
	(0.00)
Psychological Moving Cost	
Constant	6.29***
	(0.01)
Income	-0.04***
	(0.02)
Age: 35 - 54	1.01***
	(0.01)
Age: 55 +	1.93***
	(0.01)
Couple	0.13***
	(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.

Table A5 Determinants of Flow Utility $\hat{u}$

	(1)	(2)	(3)
High Quality	0.26***	0.44***	0.15
	(0.08)	(0.14)	(0.10)
Low Quality	−0.03	0.08	0.14*
	(0.06)	(0.10)	(0.08)
Three Rooms	−0.18	−0.44**	−0.05
	(0.14)	(0.22)	(0.17)
Four Rooms and Above	−0.50***	−1.08***	−0.10
	(0.07)	(0.11)	(0.08)
Home-Ownership	0.76***	0.42***	0.56***
	(0.06)	(0.10)	(0.07)
High Quality × Home-Ownership		−0.38***	
		(0.15)	
Low Quality × Home-Ownership		−0.25**	
		(0.10)	
Three Rooms × Home-Ownership		0.58***	
		(0.19)	
Four Rooms and Above × Home-Ownership		1.28***	
		(0.10)	
High Quality × Income			0.15***
			(0.06)
Low Quality × Income			−0.24***
			(0.07)
Three Rooms × Income			−0.19***
			(0.05)
Four Rooms and Above × Income			−0.52***
			(0.04)
Home-Ownership × Income			0.27***
			(0.04)
Year FE	Yes	Yes	Yes
Household Type FE	Yes	Yes	Yes
Observations	78,870	78,870	78,870
Adjusted R^2	0.67	0.74	0.68

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 A6 Willingness to Pay for House Attributes: Renters vs Owners

	Income Groups				
	1	**2**	**3**	**4**	**5**
A. Willingness to Pay (Renters)					
High Quality	61.8	74.8	90.9	114.1	180.1
Low Quality	35.4	28.6	20.2	8.1	−26.4
Three Rooms	−58.3	−72.6	−90.2	−115.6	−188.1
Four Rooms and Above	−160.4	−188.4	−222.8	−272.7	−414.6
B. Willingness to Pay (Owners)					
High Quality	1.3	5.8	11.3	19.2	41.9
Low Quality	−17.1	−25.5	−35.9	−50.8	−93.3
Three Rooms	36.6	34.1	31.0	26.5	13.7
Four Rooms and Above	60.3	52.8	43.7	30.4	−7.3

Notes: This table reports the average willingness to pay for housing attributes of households with different income levels, differentiating between renters (Panel A) and homeowners (Panel B). Income groups are defined based on yearly household income in NOK of <400k (group 1), 400-600k (group 2), 600-800k (group 3), 800-1,000k (group 4), >1,000k (group 5). In each panel, the first row shows 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 shows 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 shows how much an average household in each 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.

Table A7 Hedonic Rental Price and Investors' Demand

	log(**Rental Price**)	log(**Investors' Holding**)
Home-ownership Rate	0.37***	
	(0.03)	
Rental Yield		25.22***
		(9.38)
Product Fixed Effects	Yes	
Time Fixed Effects	Yes	
District Fixed Effects		Yes
Size Fixed Effects		Yes
Observations	360	360
Adjusted R^2	0.99	0.36
Mean Home-ownership Rate	0.64	
Std Dev Home-ownership Rate	0.18	
Mean Rental Yield		0.05
Std Dev Rental Yield		0.01

Notes: This table reports the regression models of the hedonic rental price (Column 1) and investors' demand (Column 2). The dependent variables are the logarithm of rental prices and the stock holdings of investors for each type of products between 2011 and 2018. Home-ownership Rate is the fraction of properties, for each housing product j at time t, that is owned by households and used as main residence. Rental Yield is the ratio of rental price to house price. *p<0.1; **p<0.05; ***p<0.01.

Table A8 Long-Term Effect of Change in LTI Limit on Choice Sets and Moving Probabilities

		Income Groups				
	LTI	**1**	**2**	**3**	**4**	**5**
All households						
Owners' Share of Total	Base	52.9	61.3	68.4	73.3	73.1
Products in Choice Set	Δ	−14.3%	−18.8%	−17.5%	−15.2%	−1.0%
Owners' Share of High Quality	Base	14.2	16.0	18.7	21.3	22.9
Products in Choice Set	Δ	−12.5%	−19.4%	−26.2%	−20.8%	−3.2%
Owners' Moving Probability	Base	6.7	6.8	5.8	5.6	5.4
	Δ	−1.6%	0.7%	−0.7%	−1.7%	−0.3%
Renters' Moving Probability	Base	15.6	35.6	35.5	33.2	30.7
	Δ	3.8%	−2.0%	−1.8%	−1.7%	−0.6%
Movers						
From Renting to Owning	Base	3.8	13.5	23.8	31.3	36.0
	Δ	−48.2%	−48.9%	−18.9%	−11.0%	−2.9%
From Low to High Quality	Base	5.1	6.5	7.6	9.5	10.9
	Δ	−28.2%	−9.4%	−14.8%	−12.8%	−4.1%

Notes: This table reports the longer-term (two years) 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.

Table A9 Effect of LTI Change on Aggregate Household Leverage

		Income Groups				
	LTI	**1**	**2**	**3**	**4**	**5**
All households						
Lower Bound Debt	Base	40.1	89.8	80.1	46.3	67.6
	Δ	−0.3%	−0.4%	−0.5%	−0.4%	−0.1%
Upper Bound Debt	Base	41.4	92.4	81.7	46.1	68.5
	Δ	−0.5%	−0.6%	−0.2%	−0.5%	−0.1%

Notes: This table presents the sum of total debt (in billions of NOK) of all households across five income groups for the baseline scenario (Base) and the percentage changes under the counterfactual scenario of changing LTI limits (Δ). The top two rows reports the lower bound assumption of debt, where households use all their wealth to finance their home purchase and financial moving costs. The bottom two rows reports the upper bound assumption of debt, where households use only housing wealth to finance their home purchases, and finance the rest by debt.

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:

$$A_{i,t} = P_{j,t} + \text{Financial Wealth}_{i,t} + \text{Other Real Estate}_{i,t} - \text{Debt}_{i,t}, \tag{25}$$

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, age^2, age^3, 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$\{$(Total j units/Total households in d) $\times$ Total households in d in our sample$-$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.

We define homeowners as all households with real estate wealth reported in tax returns above the median assessed tax value of dwelling by year.[34] This approach results in a homeownership rate in the sample almost identical to that reported by Statistics Norway.

In addition to the household sector, we include a real estate investors' sector that transacts in the housing market to maximize risk-adjusted profits. Households that buy or sell multiple units in a year are re-classified as investors. All housing units that are not owned by households are classified as units owned by investors.

[34]Source: https://www.ssb.no/en/statbank/table/09838/.

Appendix C Supplemental Descriptive Evidence

C.1 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.[35] In the regressions displayed in Table A10 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 A1 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 A10 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.

C.2 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-

[35] We downloaded these mortgage data from finansportalen.no.

Table A10 Determinants of Mortgage Interest Rates

	Nominal Mortgage Interest Rate (with Fees)					
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	3.12***	2.60***	4.14***	2.71***	2.76***	2.77***
	(0.14)	(0.27)	(0.16)	(0.34)	(0.29)	(0.28)
Post-LTI						-0.01
						(0.08)
Policy Rate	0.70***	1.51**		1.48**	1.47**	1.46**
	(0.04)	(0.34)		(0.34)	(0.33)	(0.33)
(Policy Rate)^2		-0.31**		-0.35**	-0.35**	-0.35**
		(0.08)		(0.08)	(0.08)	(0.09)
(Policy Rate)^3		0.03**		0.04***	0.04***	0.04***
		(0.01)		(0.00)	(0.)	(0.01)
DNB					-0.21***	-0.21***
					(0.03)	(0.03)
Handelsbanken					-0.38**	-0.38**
					(0.13)	(0.12)
Sparebanken					0.22***	0.22***
					(0.03)	(0.03)
Nordea					-0.01	-0.01
					(0.03)	(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 (≤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.

quality neighborhood, and use the middle-quality neighborhoods as the reference group (i.e., $D_{i,1990} = 0$) in both specifications:

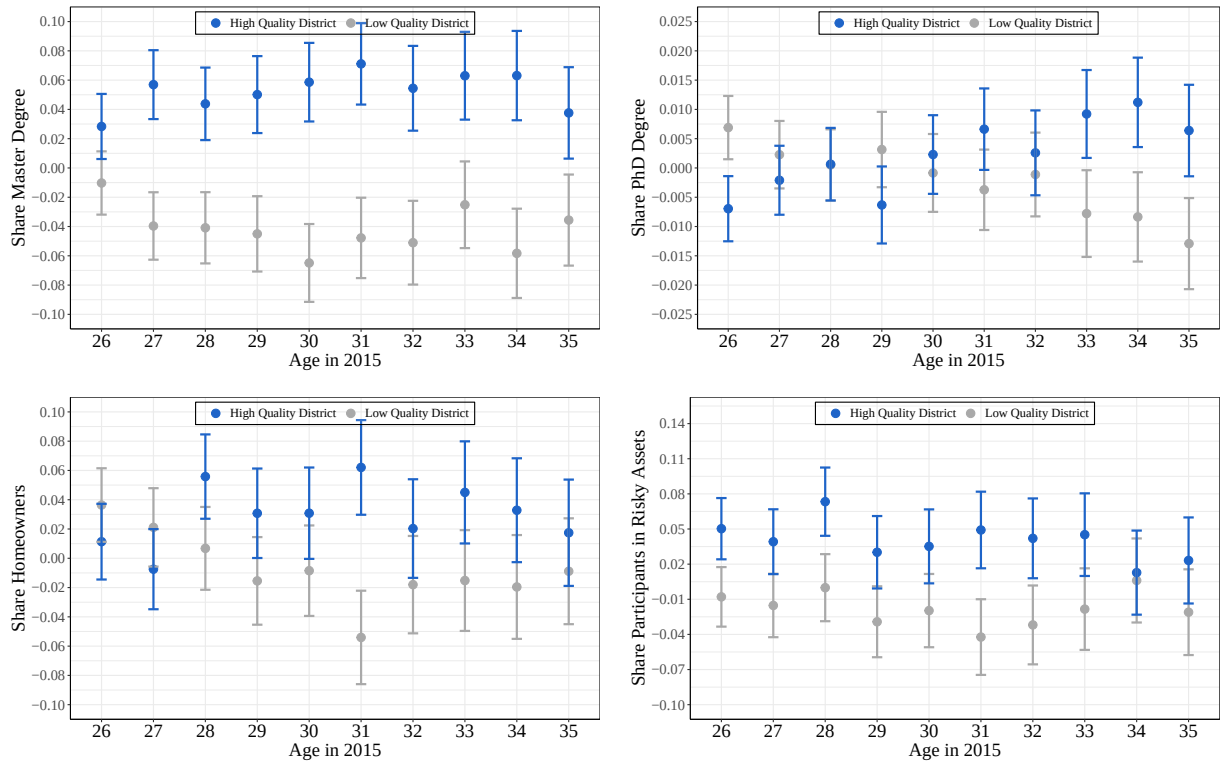
$$I_{i,2015} = \sum_{j=26}^{35} \gamma_j \mathbb{1}\{age_{i,2015} = j\} + \sum_{j=26}^{35} \eta_j \mathbb{1}\{age_{i,2015} = j\} \times D_{i,1990} + \delta E_{i,1990} + \epsilon_{i,2015}. \quad (26)$$

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 A2 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. The same figures based on regressions without controlling for parents' education show that 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 district quality plays a pivotal role in social mobility.



Notes: The figure plots the coefficients $\hat{\eta}_j$ from equation (26) with confidence intervals.

Figure A2 Outcomes when Growing up in Low vs High-Quality District

Appendix D Kernel Smoothing

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:

$$\widehat{\widehat{\Pr}}^{\bar{\tau}}_{j,h,t} = \frac{\sum_{i=1}^{N} \mathbb{1}_{[d_{i,t}=\{j,h\}]} \cdot W^{\bar{\tau}}\left(\bar{Z}_{i,t}\right)}{\sum_{i=1}^{N} W^{\bar{\tau}}\left(\bar{Z}_{i,t}\right)}, \tag{27}$$

where $W^{\bar{\tau}}\left(\bar{Z}_{i,t}\right)$ 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:

$$W^{\bar{\tau}}\left(\bar{Z}_{i,t}\right) = \prod_{l=1}^{L} \frac{1}{b_l^{\bar{\tau}}} N\left(\frac{\bar{Z}_{i,t}(l) - \bar{Z}^{\bar{\tau}}(l)}{b_l^{\bar{\tau}}}\right), \tag{28}$$

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