The London School of Economics and Political Science

Essays in Real Estate Economics and Household Finance

Sidharth Moktan

A thesis submitted to the Department of Economics of
the London School of Economics and Political Science
for the degree of Doctor of Philosophy.

London, October 2025.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 40,000 words.

Statement of Joint Work

I certify that Chapter 3 of this thesis is based on Moktan et al. (2025) and jointly authored with Benjamin Guin and Liam Clarke. I contributed 90% of the work in Chapter 3.

To Mum and Dad, my constant champions,

for always encouraging me to try, learn, and grow.

Acknowledgements

I am deeply grateful to my advisors Alessandro Gavazza and Daniel Sturm for their invaluable guidance and constant support.

A special thanks to Lin Fan, Ananya Kotia, and Ameek Singh for constant feedback and encouragement throughout the project.

I am grateful to my co-authors Benjamin Guin and Liam Clarke at the Bank of England for their co-authorship and support. The views in this thesis are my own, and not necessarily those of the Bank of England or its committees.

For helpful comments I thank Thomas Bezy, Shania Bhalotia, Julio Brandao-Roll, Gharad Bryan, Donald Davis, Gaia Dossi, Shadi Farahzadi, Michael Gmeiner, Alan Manning, Guy Michaels, Rachel Ngai, Dennis Novy, Martin Pesendorfer, Veronica Rappaport, Rebecca Rose, Mark Schankerman, Pasquale Schiraldi, Olmo Silva, Caterina Soto-Vieira, Catherine Thomas, and seminar participants at LSE.

Abstract

This thesis consists of three essays on real estate and household finance.

The first chapter examines heterogeneity in housing markets both within and across the largest cities in the UK. I combine several novel datasets of UK property markets and document three key novel facts: (i) housing quality is segmented between the rental and owner-occupied sectors, with rentals generally offering lower quality, (ii) cities with more quality segmentation tend to have higher rent-to-price ratios, and (iii) in more segmented cities landlords have fewer assets and face lower capital gains. These facts suggest that heterogeneity in landlord supply constraints may play an important role in shaping local housing markets.

In the second chapter, I develop and estimate a two-sided assignment model of the housing market to quantify the effect of landlord supply constraints. The model features households' optimal choice of housing quality and tenure (i.e., the choice to rent or own) in the presence of borrowing constraints, landlords' profit-maximizing choice of quality to rent out, and endogenous quality segmentation and rent-to-price ratios determined in equilibrium. I conduct counterfactual experiments to show that differences in landlord supply constraints explain much of the variation in quality segmentation and rent-to-price ratios observed within and across cities.

The third chapter, based on work co-authored with Benjamin Guin and Liam Clarke, explores the role of lending in transmitting shocks to residential real estate. We examine an adverse and salient shock to a segment of the property market (high rise flats) in England and Wales. Using an event-study design that compares flats and non-flats in the same local area, we show that mortgage originations for affected

properties declined sharply by around 30% in London following the change. This contraction was concentrated among first-time buyers who experienced an almost 50% drop. At the same time, the share of cash purchases rose by five percentage points, implying a re-allocation of ownership from credit-constrained households to wealthier, mortgage-independent buyers. To distinguish between reduced credit supply and falling borrower demand, we examine the evolution of rents and prices. While prices fell by about 2.5% relative to unaffected units, rents rose by 2%, suggesting that demand to reside in these homes remained stable even as their collateral value declined.

Contents

List of Figures

List of Tables

Chapter 1

Heterogeneity in Housing Markets Within and Across Cities*

Homeownership rates, prices, and rents vary dramatically both within and across local housing markets. Differences in the availability and affordability of housing have important implications for the distribution of household welfare because housing is both a major source of expenditure for households and a key determinant of access to local amenities such as schools. Understanding these differences is particularly timely given the recent rise in the private rental sector as shown in Figure 1. Despite its importance to households, we know little about why these differences arise, how they affect the distribution of household welfare across cities, and what their implications are for housing policy.

I answer these questions in the first two chapters. I start out in the first chapter by documenting several new facts which highlight differences in housing markets within and across England's largest cities and point to the potential role of landlords in generating these patterns.[1] To illustrate these patterns, I compile data

*Chapters 1 and 2 use Crown-copyright statistical data from the Office for National Statistics (ONS) and the Consumer Data Research Centre (project ID CDRC 1015-01, ES/L011840/1; ES/L011891/1.). These chapters also use data available under Open Government Licence v3.0. from HM Land Registry, the Department for Levelling Up, Housing & Communities (English Housing Survey), and the Valuation Office Agency. The use of the data does not imply the endorsement of the data provider in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

[1] In the second chapter I estimate an equilibrium model of housing markets to study the role of landlords' supply of rental properties in shaping market outcomes and generating the observed

Figure 1: Tenure Shares Over Time in London for Households Ages 20-39

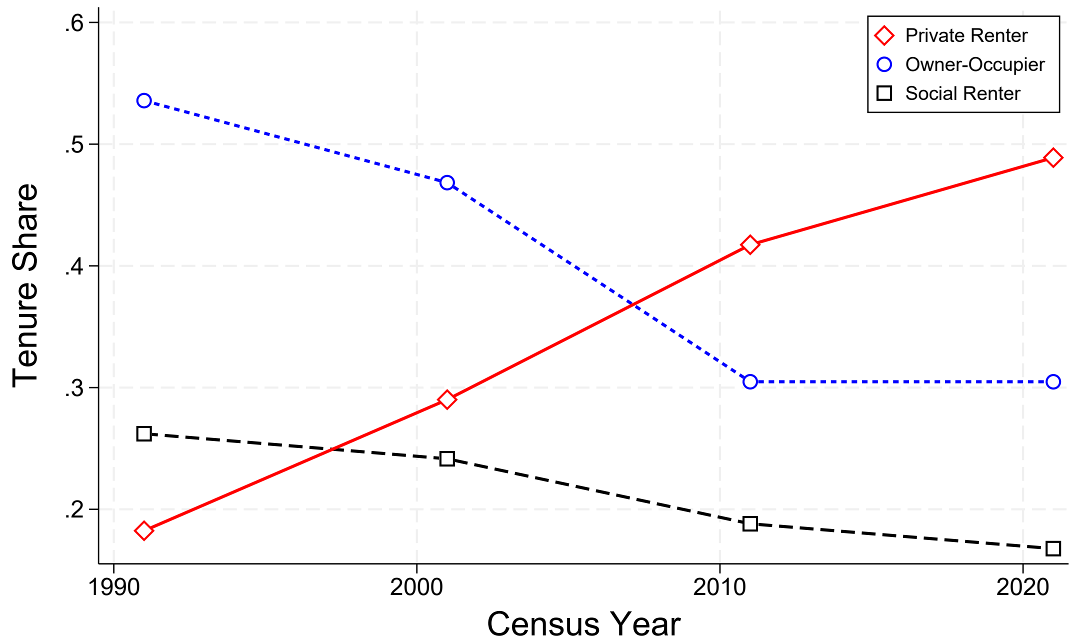

Note: This figure plots the share of households between age 20-39 who belong to one of three housing tenure categories over time. Appendix figures A.1 and A.2 present analogous time series of tenure shares for age groups and regions not plotted in this figure.

from Zoopla—the UK's second-largest property listings platform—and augment it with administrative records on property transactions from the Land Registry Price Paid dataset as well as household survey data from multiple sources.

I first examine how housing quality differs between the rental and owner-occupied sectors within different cities. I develop a method for measuring differences in the distribution of house quality between the rental and owner-occupied sectors that relies on comparing the rank achieved by the same property in the local rent and price distributions. This approach is closely related to Diamond and Diamond

systematic differences in these outcomes.

(2023, 2024) who use similar rank-rank comparisons to impute rents for owner-occupied properties in the US. A key benefit of this approach is that it decouples the measurement of relative quality by tenure which depends on ranks from the measurement of the rent-to-price ratio which depends on levels. This allows me to study how differences in rents and prices impact the distribution of quality by tenure. To implement this comparison of quality between the rental and owner-occupied sectors, I exploit data on properties that were listed both for rent and for sale in Zoopla.

These comparisons yield two insights. First, I find that housing quality tends to be segmented between the rental and owner-occupied sectors, with the rental sector drawing more heavily from the lower quality segments of the local housing distribution. Second, the degree of segmentation varies widely across cities. For instance, while the quality of the housing stock is comparable between the rental and owner-occupied sectors in London, in Birmingham, the median-quality rental has lower quality than 80% of owner-occupied properties. This highlights stark disparities in the quality of rental housing relative to owner-occupied housing across cities.

To examine why the degree of quality segmentation varies across cities, I explore differences in rents and prices across cities. I find that cities with greater quality segmentation tend to have higher rent-to-price ratios. This positive correlation suggests that differences in rental supply may play a role in generating the observed patterns of segmentation. In cities like London renters may choose to live in relatively better houses because landlords are willing to accept lower rents relative to prices.

The positive correlation also implies that the rent-to-price ratio does not solely reflect the relative quality of rentals compared to owner-occupied properties. If it did, the rent-to-price ratio would be higher in cities in which renters live in relatively better housing, generating a negative correlation between the rent-to-price ratio and the degree of segmentation (Head et al., 2023). The presence of a positive correlation suggests that the rent-to-price may instead be driven by differences in supply constraints faced by landlords across different cities.

To explore whether the observed market outcomes are consistent with a mechanism in which landlord supply constraints drive differences in the rent-to-price ratio, I investigate the relationship between the aggregate rent-to-price ratio in a city and several city-level factors which influence aggregate landlord supply constraints in a city. First, I show that the aggregate rent-to-price ratio tends to be lower in cities in which there is greater aggregate levels of capital gains, measured using repeat sales data for houses sold both in 2010 and 2018. Second, the higher prices in low rent-to-price cities such as London imply that land accounts for a larger share of property values, suggesting lower effective depreciation rates. Third, recent evidence by Amaral et al. (2025) indicates that housing assets in large cities such as London are perceived as safer and more liquid, implying lower risk premia. Taken together these patterns suggest that quality might be less segmented in cities like London because landlords in these cities are willing to supply rental units for low rents relative to prices given that they face lower user costs net of capital gains.

One concern with this interpretation is that landlords may rely more heavily on borrowing in cities such as London due to higher property prices. If this were the case, higher borrowing costs might offset the user-cost advantages in London stemming from lower depreciation, lower perceived risk, and higher expected capital gains. Using data from the Wealth and Assets Survey, I compare loan-to-value ratios on buy-to-let properties between landlords in London and those in the rest of England, and find that the loan-to-value ratios in London are not statistically significantly different from those elsewhere. To investigate why the loan-to-value ratio does not vary despite large differences in prices, I examine differences in landlord wealth across cities. I find that in cities such as London, the wealth gap between landlords and households is much larger than in cities with higher rent-to-price ratios. This greater wealth allows landlords in London to maintain similar leverage levels despite much higher property prices.

While the aggregate rent-to-price ratio is strongly decreasing in aggregate capital gains across cities, I show that when comparing across quality segments within a city the correlation either disappears or reverses in sign and becomes positive. This implies that the within-city variation in the rent-to-price ratio cannot be explained by offsetting differences in capital gains and may instead arise due to variation in other components of landlord user costs. I show in Chapter 2 that variation in landlord borrowing levels and borrowing costs can explain almost all of the observed variation in rent-to-price ratios within London.

1.1 Data

I combine information from different datasets on houses and households. To estimate distributions of prices and rents by quality of housing, I use a comprehensive dataset of rental and sale listings from Zoopla, the second-largest real estate listings platform in the UK. To measure households' characteristics and housing choices, I use data from two different household surveys. The combined data is used for both Chapters 1 and 2.

Listings and Transactions—I obtain listing-level data from the WhenFresh-Zoopla dataset. The dataset records the universe of rental and sale listings in Zoopla, the second-largest listing platform in the UK. The dataset covers the period 2014–2020 and includes over five million listings of properties for sale and over four million properties for rent over this period. For each listing, the data records the listing date, transaction date, prices for properties on sale, weekly rents for rentals, and property characteristics including property address, number of rooms, and build type.

Household Choices, Endowments, and Characteristics—I use data from the Wealth and Assets Survey (WAS) and the UK Household Longitudinal Survey (UKHLS). The WAS is a biennial longitudinal survey of household assets and savings. During the period 2014-2016, the WAS includes approximately 18,000 households per wave or approximately $9,000$ households per year. I use the WAS to estimate joint distributions of wealth and income by region, year, and household characteristics for households. I also use it estimate the asset distribution for landlords by region and year.

The UKHLS is the largest longitudinal household survey in the UK. Over the period 2014-2016, the survey includes approximately $22,000 - 25,000$ households per year. I use the UKHLS to measure the relationship between households' characteristics and housing choices by city and year. It includes ownership rates by household age, correlations between household income and rent for renters, and correlations between household income and house price for owner-occupiers.

Other Datasets—I combine data on the housing stock from the 2011 Census and Valuation Office Agency (VOA) to construct weights which are used to estimate the distribution of prices and rents over the entire housing stock. I use the Census to measure time trends in ownership rates over a long time horizon. I also use administrative records of property transactions from the Land Registry Price Paid Data to estimate capital gains in the housing market.

1.2 Comparing the Distribution of Quality Between the Rental and Owner-Occupied Sectors

In this section, I describe how I measure the distribution of prices and rents over quality as well as differences in the distribution of quality between the rental and owner-occupied sectors. Following the literature on hedonic models and models of housing assignment (Landvoigt et al., 2015; Epple et al., 2020; Diamond and Diamond, 2023, 2024), I represent housing quality using a unidimensional index which captures all relevant features of a house and its environment. Prices and rents are increasing in this quality index. The monotonic relationship between rents and quality implies that I can use a property's rent rank in any given city and year to recover its quality rank in the rental sector in that city and year. Similarly, a property's price rank can be used to recover its quality rank in the overall housing market of a given city during a given year.

However, several complications arise when trying to compare the distribution of quality, prices and rents over time and across the rental and owner-occupied sectors. First, while I need distributions over the entire stock of housing, I only observe prices or rents for listed properties. To recover distributions over the entire stock, I need to adjust for over- or under-representation of properties in the listings data. Second, the price or rent rank for a given level of quality may change over time as properties are either added or removed from the stock. To recover the path of prices or rents for a given unobserved level of quality, I need a way to trace the ranking of this quality level in the price and rent distributions over time. Third, comparing quality across the rental and owner-occupied sectors is complicated because I do not observe a common measure of quality for properties in the two sectors. In general prices are not observed for rental listings and rents are not observed for sale listings. To resolve this missing data problem, I need a strategy to map the rental quality distribution to the quality distribution of the overall housing stock.

In this section I discuss how I resolve each of these measurement issues to obtain quality distributions which are comparable over time and across the rental and owner-occupied sectors. These adjusted distributions form the basis for how I measure the evolution of rents and prices for different levels of unobserved quality. Using these adjusted distributions, I find novel evidence that the quality of the housing stock tends to be lower in the rental sector compared to the owner-occupied sector. I also show that this documented pattern of quality segmentation by tenure varies systematically across cities.

1.2.1 Measuring the distribution of quality, rents, and prices

Marginal Distributions of Rents and Prices—We are interested in the overall distribution of rents and prices at the city and year level. However, the data only lets us observe prices and rents for listed properties. Since listed properties might not be representative of the stock of houses in the rental sector and in the housing market as a whole, the observed distributions need to be adjusted for selection into listing. To obtain the adjusted distributions, I first measure the degree to which each type of property is either over- or under-represented in the listings data relative to the overall stock. I then use this information to estimate weighted distributions of rents and prices which are representative of the stock of rented properties as well as the entire housing stock.

For this exercise, I define a property type as a combination of Lower Layer Super Output Areas (LSOA) and number of bedrooms.[2] I choose the LSOA as the spatial unit because the Valuation Office Agency's (VOA) data on the housing stock is provided at this level of granularity. Let P^r_{jcy} and P^s_{jcy} represent the share of type j properties among the rental stock and overall housing stock. I construct these shares by combining yearly data from the VOA with the 2011 census. Let p^r_{jcy} and p^s_{jcy} represent the share of type j properties among rental and sale listings in city c in year y. Then p^r_{jcy}/P^r_{jcy} and p^s_{jcy}/P^S_{jcy} measure the degree to which properties of type j are

[2]There are 33,755 LSOAs in England each comprising 400-1,200 households with a resident population between 1,000-3,000 persons.

over or under-represented among rental and sale listings respectively. Weighting the rental and sale listings by the inverse of these ratios, I recover the distribution of rents in the rental sector $F_r(\cdot \mid city, year)$ and the distribution of prices over the housing stock $F_p(\cdot \mid city, year)$ by city and year.[3]

The weighting strategy assumes that the selection of properties into the sample of Zoopla listings only depends on observable attributes of the property. This assumption will be violated if, after conditioning on observable property attributes, the distribution of unobserved quality among listed properties differs from that in the overall housing stock. For example, if listed properties tend to have worse unobserved quality, then the distribution of rents, $F_r(\cdot \mid city, year)$, and prices, $F_p(\cdot \mid city, year)$, estimated from the Zoopla listings will be underestimated relative to the true distribution for the housing stock. The assumption can also be violated if sellers or landlords who list on Zoopla differ from the overall population of potential sellers or landlords respectively. For instance, landlords who list on Zoopla may be more active than the overall population of landlords because they may face lower frictions of re-listing properties or may be more aware of market conditions such as increases in market rents. Such differences may bias upwards the distribution of rents estimated from the Zoopla listings.

Normalizing House Quality—Quality needs to be normalized because it is unobserved. In each city, I define a property's relative quality within the city by its rank in the city's overall distribution of house prices as of 2014.[4] This is equivalent

[3]To clarify, these are weighted distributions of listing-level prices or rents, which have been weighted to correct for over- or under-representation of different property types in the Zoopla dataset.

[4]Note that this definition of quality allows me to explore within-city differences in quality distributions by tenure. It can also be used to compare the relative quality by tenure across cities. However,

to normalizing quality to equal prices as of 2014 since price ranks are monotonic in prices.[5] Under this definition of quality, we only directly observe a property's relative quality if we know its price as of 2014. We need a way to infer quality for properties for which we do not observe this information. These include properties for which we only observe prices from other years as well as properties with data on rents. I develop a strategy which allows me to assign quality to such properties based on the rent or price information that is available.

Tracing the Rank of Quality Over Time Within the Rent and Price Distributions—We need prices and rents between 2014 and 2018 for each level of unobserved quality. We have already estimated the marginal distributions of prices and rents by city and year. Therefore, to track prices and rents by quality over time, we need to trace the rank of each quality level in the price and rent distributions over time.

As discussed above, we have already estimated the distributions of prices $F_p(\cdot \mid c, y)$ over the housing stock and rents $F_r(\cdot \mid c, y)$ in the rental sector by city and year. The rank of any for-sale listing l in city c and year y with price p_{lcy} is given by $F_p(p_{lcy} \mid c, y)$. Rental listings can similarly be ranked. For each tenure and city, we would ideally estimate the evolution of ranks between two years by comparing the rank attained by the same property at these two points in time. I do not use a repeat listings analysis for this exercise because the sample of repeat listings becomes very sparse for some city-year-tenure combinations, especially in smaller cities. Given the data restriction, I instead collapse the listing-level data on price and rent ranks to

I cannot directly use them to compare quality distributions across cities. In the structural estimation, I treat the distributions of quality as unobserved and estimate them.

[5] Landvoigt et al. (2015); Epple et al. (2020) normalize house quality to equal prices in a base year.

a granular property-type level by taking the median of the ranks for each property type by city, tenure and year. I then estimate the evolution of ranks using this data on median ranks. I describe this procedure in more detail below.

I define each property type i by a combination of postcode, dwelling type, and the number of bedrooms and bathrooms. The high level of granularity of this property-type definition helps mitigate concerns about aggregation. Let q^p_{icy} be the median rank achieved by all for-sale listings for properties of type i in city c and year y.[6] I assume the average quality of a type-i property remains unchanged during the period under study, i.e., $h_{icy} = h_{ic2014}$. Under this assumption, we can trace the price ranking of quality h_{ic2014} over time by tracing the ranking of properties of type i. Similarly, we can use q^r_{icy} , the median rent rank for properties of type i, to trace the ranking of quality h_{ic2014} in the rent distribution.

We need to estimate the mapping $q^p_{icy} = m_{pcy}(q^p_{ic2014})$ between initial median price rank q^p_{ic2014} and the median price rank q^p_{icy} in subsequent years $y > 2014$. We also need analogous quantile maps for the rental sector m_{rcy}. I estimate m_{pcy} for each city and year using quantile regressions of q^p_{icy} on a polynomial of q^p_{ic2014}. I obtain the quantile maps in the rental sector m_{rcy} using an analogous procedure. I use k-fold cross-validation with five folds ($k = 5$) to select the polynomial degree for each mapping m_{pcy} and m_{rcy}. Appendix Table A.2 reports the optimal polynomial degree selected using cross-validation.

[6] If a property of type i is listed only once for sale in city c in year y, then q^p_{icy} is simply the rank of this listing. However, in practice, there might be more than one for-sale listing of properties of type i in a given year within a city. In these cases, q^p_{icy} is the median across these listings. Formally, $q^p_{icy} = \text{median}\left\{F_p(p_{lcy} \mid c, y) \mid l \in \mathcal{L}^p_{icy}\right\}$ where $\mathcal{L}^p_{icy}$ is the set of all for-sale listings for properties of type i in city c and year y.

The estimation for mapping m_{pcy} uses data for property type i only if at least one property of type i was listed for sale both in 2014 and y. The estimated mappings will be biased if the evolution of price ranks for property types included in the estimation sample differ from those that are excluded. This might occur if properties that are listed for sale in both 2014 and y are more likely to have undergone renovations and improvements in the intervening period. In this case, the mapping of 2014 ranks to future ranks would be biased upwards. A similar issue exists for the estimation of mapping m_{rcy}. Given the absence of data on renovations, improvements, and landlords, I assume that the unobserved quality of a property type is fixed over the sample period. Bias can also be introduced if households or landlords who list repeatedly are different than those who do not. For instance, landlords who list the same property every year may increase rents more aggressively than those who list less frequently. This type of selection can bias the mapping of 2014 ranks to future ranks. I assume that such selection does not exist.

Appendix Figure A.3 plots the functions m_{pcy} and m_{rcy} for $y = 2018$ in Birmingham, London, and Manchester. The main patterns discussed below also hold for cities and years not plotted in the figure. There are two key takeaways. First, the plots do not coincide with the diagonal which implies that the rank of a property type or quality level is not fixed over time. We therefore need a way to keep track of the rank for each quality level over time. Second, the rank in 2018 is increasing in the rank in 2014. The monotonicity implies that for any price rank q in the 2014 distribution, we can assign a unique rank $m_{pcy}(q)$ in the year-y distribution and vice-versa. Therefore, for any quality level with initial price rank q, we can trace its ranking over time by evaluating the functions $m_{pcy}(q)$ for each year between 2015 and 2018. Moreover, for properties with price rank q in 2014, we can recover its year-y price since it equals quantile $m_{pcy}(q)$ of the estimated price distribution for year y. Repeating this for all $y > 2014$ yields the path of prices for quality levels with initial price rank q. Similarly, we can obtain the path of rents given an initial rent rank.

Mapping Quality Between the Rental and Sale Distributions—The preceding analysis lets us trace the path of prices given a price rank from 2014 and the path of rents given a rent rank from 2014. Since quality has been normalized to equal the price rank as of 2014, this is sufficient to track the path of prices for each level of the normalized quality. However, to track the path of rents by each level of normalized quality, we need to establish a mapping between the rent rank and the 2014 price rank. To estimate this link, I exploit information from properties which have been cross-listed for sale and rent in the dataset.

The estimation is conducted in two steps. In the first step, restricting the sample to properties that have been listed both for rent and sale at some point between 2014 and 2018, I assign each property ranks in the 2014 rent distribution and the 2014 house price distribution. For sales listings in 2014, we directly observe its sales rank as of 2014. For sales listings from future years $y > 2014$, I invert the mapping $m_{pcy}(\cdot)$ estimated above to assign a price rank in 2014 given a price rank as of y. I assign 2014 rent ranks to listings using an analogous procedure. To deal with properties that have been listed for sale more than once, I define a property's rank in the price distribution as the median 2014 price rank achieved by the property across all of its sale listings. Each property's rent rank is defined analogously. This procedure yields a pair of price and rent ranks as of 2014 (q_{p2014}, q_{r2014}) for each property in the restricted sample.

In the second step, I use the assigned ranks to estimate a mapping $g_c(q)$ which predicts a property's rank in the 2014 price distribution given its rank in the 2014 rent distribution. For each city, $g_c(q)$ is estimated using a quantile regression of each property's 2014 price rank against a polynomial of its 2014 rent rank. I select the degree of the polynomial for the quantile regression in each city using a k-fold cross-validation with five folds ($k = 5$). Appendix Table A.1 reports the polynomial order used to estimate $g_c(q)$ by city. The regression for each city is estimated over all properties in the restricted sample.

The function $g_c(q)$ links the ranks of the 2014 rent distribution to the ranks of the 2014 price distribution. Since we normalize quality to equal the 2014 price ranks, $g_c(q)$ also provides a link between the rent distribution and the distribution of quality. This mapping forms the basis for how I impute prices for rental properties and measure quality segmentation across the rental and owner-occupied markets.

The sample used to estimate $g_c(q)$ includes only those properties that were listed for rent and for sale at least once between 2014 and 2018. The estimated $g_c(q)$ could be biased if the rent rank-to-sale rank relationship for properties in the estimation sample differs from the relationship for properties excluded from the sample. Similar to the discussion of potential bias in the estimation of m_{pcy} and m_{rcy} above, this might occur if properties undergo improvements or renovations before being re-listed either for rent or for sale. The presence of such bias cannot be ruled out using the available data. I assume that the quality of properties is fixed between

2014-2018 for the analysis.

Imputing Prices for Rental Properties—Consider a rental property which lists in the q^{th} quantile of the rental distribution in city c and year y. We are interested in the price of this rental unit. However prices are generally not available for rental listings. To impute the unobserved price, I apply to q the composition of mappings $m_{pcy} \circ g_c \circ m_{rcy}^{-1}$. $m_{rcy}^{-1}(q)$ converts the rent rank from year y to the rent rank as of 2014. Applying g_c to the 2014 rent rank gives the 2014 price rank which is also the measure of quality. Lastly applying m_{pcy} to the 2014 price rank gives the price rank as of y for this property. Given the price rank as of y, we can lookup the corresponding price as of y using the price CDF. Therefore, by applying the appropriate composition of mappings, we can impute prices in any year between 2014 and 2018 given a rent rank in this time period. Dividing the rent by the imputed price yields the imputed rent-to-price ratio for a given level of rent.

1.2.2 Quality Segmentation

Quality is Segmented by Tenure. Segmentation Varies by City—I use the estimated mapping $g_c(q)$ to infer the relative quality of properties in the rental and owner-occupied sectors. To make inferences about relative quality based on $g_c(q)$, prices need to be monotonically increasing in rents. As discussed in the following paragraph, I verify that this condition is satisfied in the data. When this condition is satisfied, $g_c(q) < q$ implies that quality in quantile q of the rental distribution is lower than quality in the same location of the owner-occupied distribution. $g_c(q) > q$ implies the opposite.

Figure 2: Rent Rank vs. Predicted House Price Rank

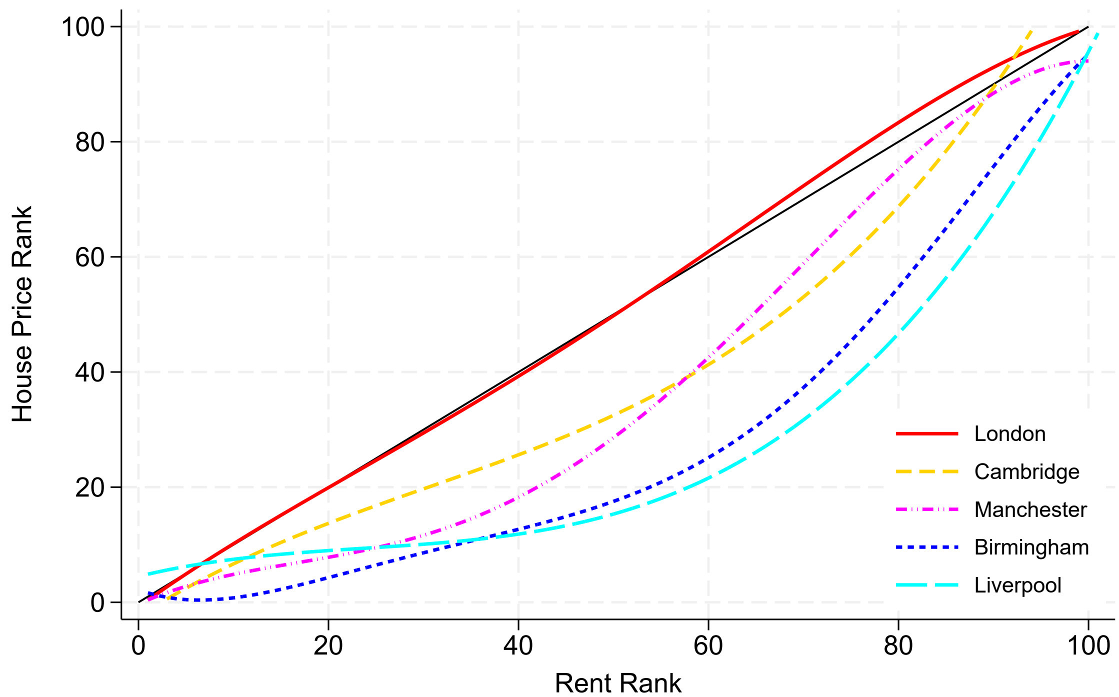

Note: This figure plots the conditional rank in the overall distribution of house prices given a property's rank in the overall rent distribution. Each plot corresponds to a different city.

$g_c(q) = q$ arises as a special case when quantile q of both distributions have the same quality.

Figure 2 plots the mapping $g_c(q)$ for five cities. Three key patterns emerge. First, the price rank $g_c(q)$ is monotonically increasing in the rent rank q across all cities. This verifies that prices are monotonically increasing in rents. Second, $g_c(q)$ is almost universally below the diagonal across all cities. This implies that compared to the owner-occupied sector, the stock of properties in the rental sector draws more heavily from the lower quality segments of the housing distribution. Third, there is substantial heterogeneity across cities in the degree to which the rank-rank plots deviate from the diagonal. Whereas rental and sale listings in London have the same median quality, the median quality for rentals in Liverpool ranks below the 20^{th} percentile of the overall housing distribution in Liverpool. This implies that the relative quality of properties in the rental and owner-occupied sectors vary greatly across cities.

These three insights provide novel evidence that quality is segmented between the rental and owner-occupied sectors and that the degree of segmentation varies substantially across cities. Quality segmentation can arise due to several reasons. It can arise if preferences for owner-occupation are correlated with households' taste for quality or their income and asset endowments, both of which influence demand for quality. Segmentation can also arise if the costs of owner-occupation are different than rents for the same quality. The quantitative analysis allows for each of these channels.

What do the Patterns of Segmentation Reflect?—The patterns of segmentation along the unidimensional index in Figure 2 are driven by segmentation in quality along three sub-dimensions: the location of a dwelling, its structural features, and its upkeep. The data shows that the degree of segmentation along these three dimensions vary across cities. Moreover, segmentation along these dimensions are correlated across cities. The higher the relative quality of the location of rented properties compared to owner-occupied properties in a city, the higher tends to be the relative quality of structural features of rented properties. Similarly, the upkeep of rental dwellings is correlated with the other two sub-dimensions. The positive correlation suggests that factors which improve the relative quality of rented dwellings in a city do so by improving all three sub-dimensions of quality. Figure

Figure 3: Ratio of Share of Dwelling Type in Rented and Owner-Occupied Dwellings

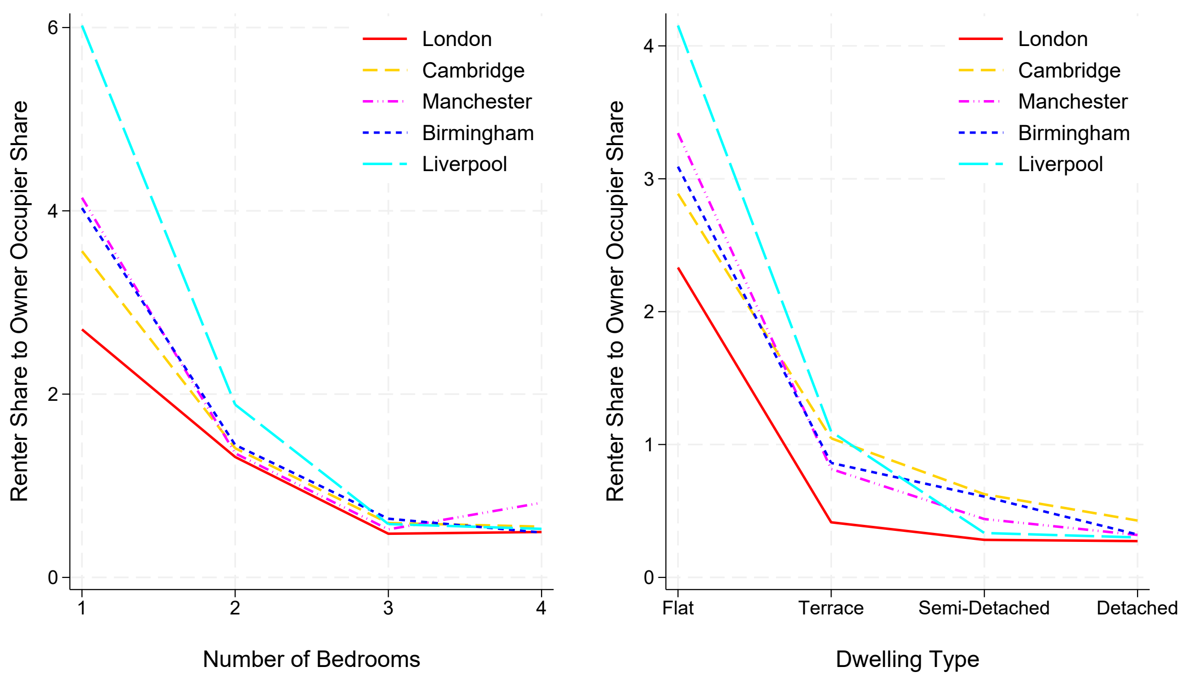

Note: This figure compares dwelling characteristics between the rental and owner-occupied sector by city. The left figure plots the share of rental units of a given size divided by the share of owner-occupied units of the same size. The figure on the right conducts the same comparison for build type.

3 shows differences in structural features of rented and owner-occupied dwellings across cities. The figure on the left plots the share of rented dwellings of a given size divided by the share of owner-occupied dwellings of the same size. The figure on the right plots the relative shares of properties with different builds. There are two key patterns to note. First, rented dwellings are smaller and more likely to be flats in every city. Rented dwellings thus tend to have structural features that command

lower value in the market. Second, the relative quality of structural features of rented dwellings tends to be lower in cities which are more segmented along the unidimensional index of quality. These patterns reveal that the segmentation along the unidimensional index is at least partly driven by segmentation with respect to the quality of dwellings' structural features.

Figure 4: Segmentation of Location Quality Between the Rental and Owner-Occupied Sectors

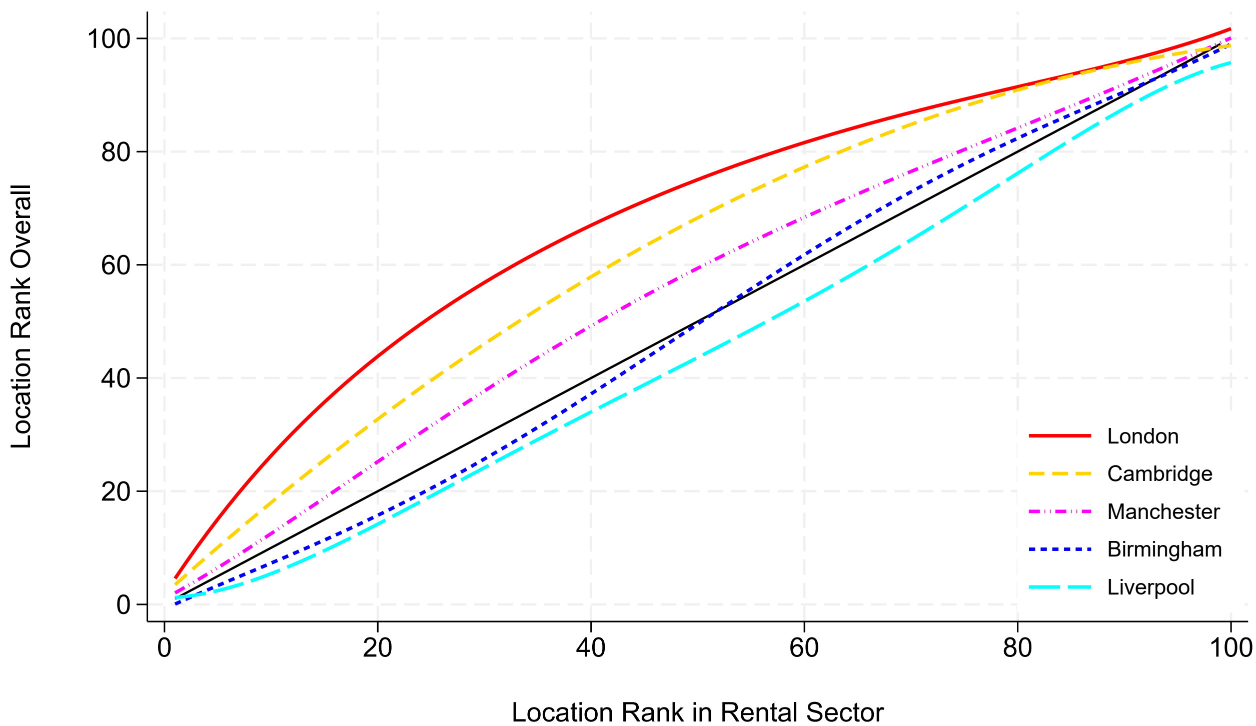

Note: This figure plots the conditional rank of a property's location in the overall distribution of location fixed effects given its rank in the distribution of location fixed effects for the rental sector. Each plot corresponds to a different city.

To measure segmentation with respect to the quality of dwellings' location, I decompose the unidimensional index of quality into a location component and a component capturing the quality of dwellings' structural features. I estimate the following fixed effects regression using data on sale listings:

$$price_{ilcy} = \alpha_y^c + \alpha_l^c + \alpha_{build(i)}^c + \varepsilon_{ilcy}$$

where i indexes property type; l indexes the LSOA; α_y^c is a fixed effect for year y; α_l^c is a fixed effect for LSOA l; and $\alpha_{build(i)}^c$ is a vector of dwelling structure fixed effects obtained by interacting the type of each dwelling with the number of bedrooms and bathrooms. The regression allows me to isolate the location component of quality in α_l^c and the structural features component in $\alpha_{build(i)}^c$. A higher value of α_l^c implies a higher quality for location l. Similarly, a higher value for $\alpha_{build(i)}^c$ implies a higher quality for dwelling feature of type $build(i)$. Using this decomposition, I construct year- and city-specific distributions of location quality in the rental sector and in the housing stock as a whole. Based on these distributions, Figure 4 plots the conditional rank of location quality in the overall housing distribution given its rank in the rental sector.[7] The figure shows that in most cities rental properties tend to be located in higher quality LSOAs compared to owner-occupied properties. Moreover, the relative quality of LSOAs of rented properties is higher in cities in which rental properties have higher relative values in the unidimensional index of quality.

Lastly, Figure 5 plots the degree of segmentation in a city against: (i) the share of non-decent dwellings in the private rental sector, and (ii) the ratio of the share of non-decent dwellings in the private rental and owner-occupied sectors. I assign

[7] The rank-rank relationship is estimated using a quantile regression of the overall rank against a polynomial of the rank in the rental sector. The polynomial order, selected for each city based on five-fold cross-validation, is two for Cambridge and five for the other cities.

Figure 5: Prevalence of Non-Decent Rental Dwellings Correlated with Segmentation

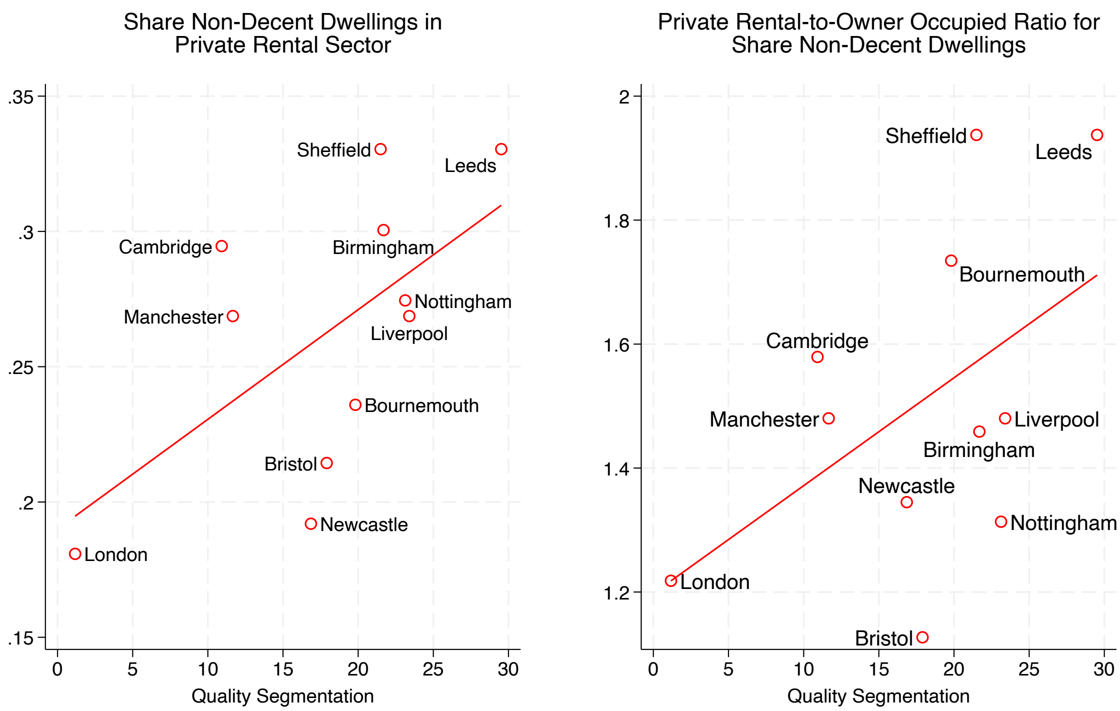

Note: This figure plots the relationship between the degree of segmentation in a city against: (i) the share of non-decent dwellings in the private rental sector; and (ii) the ratio of the share of non-decent dwellings in the private rental and owner-occupied sectors.

shares of non-decent dwellings to each city using the English Housing Survey.[8] The survey categorizes dwellings as decent or non-decent based on the Decent Homes Standard which assesses hazards, state of repair, modern facilities, and thermal comfort.[9] Figure 5 shows that in more segmented cities the condition of rental

[8]The English Housing Survey includes location information at the Government Office Region (GOR) level. I proxy the share of non-decent dwellings in each city using the share in the corresponding GOR.

[9]The Decent Homes Standard was introduced by the UK government in 2000 and is overseen by the Department for Levelling Up, Housing and Communities (DLUHC). It sets the criteria for housing quality with respect to safety, condition, amenities, and energy efficiency.

dwellings tends to be lower both in absolute terms and relative to the condition of owner-occupied dwellings.

Taken together, Figures 3, 4 and 5 show that the three dimensions of quality tend to move in tandem across cities. This suggests that trade-offs along these margins are not key to generating the patterns of segmentation observed across cities. Instead factors which improve relative quality in the rental sector of a city do so by simultaneously improving the structural features, location, and upkeep of rented dwellings.

1.3 Systematic Variation in Housing Market Outcomes and Landlord Supply Constraints

1.3.1 Cities with Higher Relative Quality of Rentals Have Lower Mean Rent-to-Price Ratios

Several studies document that the rent-to-price varies both within and across cities (Himmelberg et al., 2005; Gyourko et al., 2013; Bracke, 2015; Halket et al., 2020; Head et al., 2023). I present new evidence showing that the mean rent-to-price ratio in a city is decreasing in the relative quality of rentals in the city.

I define the degree of quality segmentation between the rental and owner-occupied sectors of a city as $\gamma_c = \sum_q (g_c(q) - q)^2$, the sum of squared differences between the rent ranks and the predicted price ranks. A higher value of γ_c implies greater segmentation.

Figure 6: Cities with Higher Relative Quality of Rentals Have Lower Mean Rent-to-Price Ratios

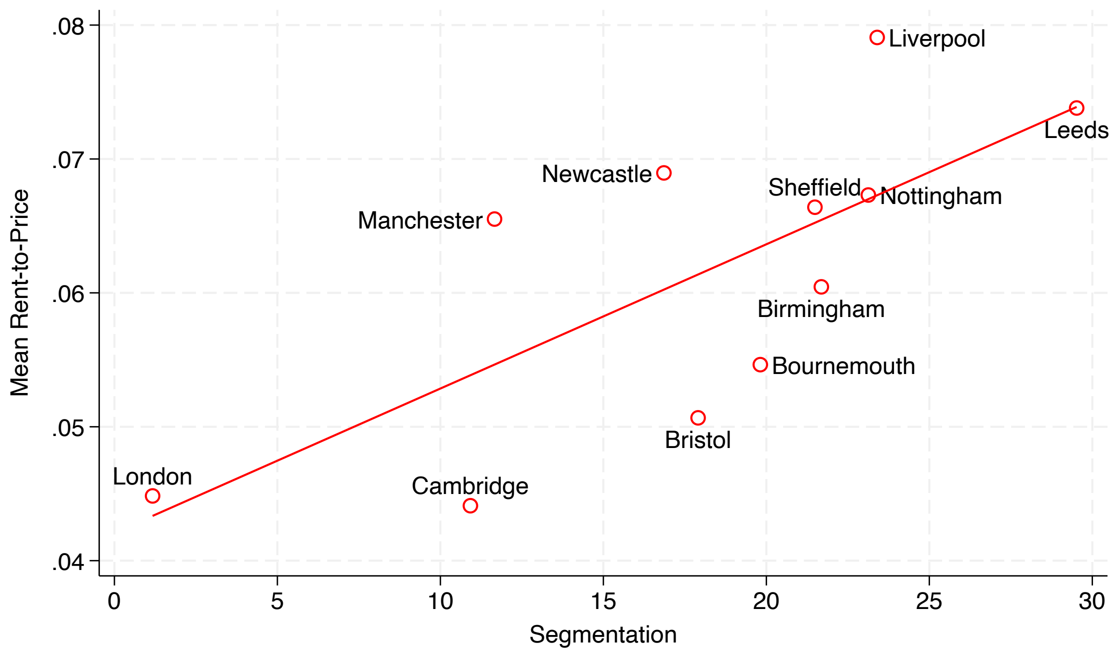

Note: This figure presents a scatter plot of the mean rent-to-price ratio in a city against the degree of quality segmentation between the rental and owner-occupied sectors in the city. The figure also plots a line of best fit.

To obtain the mean rent-to-price ratio, I rely on the mappings estimated in Section 1.2.1. Using these mappings I impute prices for each rent percentile of the rental stock by city and year. Taking the ratio of rents and imputed prices yields an imputed rent-to-price ratio for each percentile of the rent distribution. Let $\widehat{RTP}_{qcy}$ be the imputed rent-to-price ratio for percentile rank q of the rent distribution. For each city and year, the mean rent-to-price ratio for properties that belong to the rental stock is estimated as $\overline{RTP_{cy}} = \frac{1}{100}\sum_{q=1}^{100} \widehat{RTP}_{qcy}$. The average $\overline{RTP_{cy}}$ weights each percentile rank q of the rent distribution equally because the weights for the pecentile ranks in the rent distribution are uniform by construction. I summarize the mean rent-to-price in each city by averaging $\overline{RTP_{cy}}$ over 2014-2018: $\overline{RTP_c} = \frac{1}{5}\sum_{y=2014}^{2018} \overline{RTP_{cy}}$.

Figure 6 plots the relationship between the degree of segmentation, γ_c, and the mean rent-to-price ratio in each city over 2014-2018, $\overline{RTP_c}$. The plot shows that segmentation is increasing in the average rent-to-price in a city. In other words, the cost of renting relative to prices is lower in cities with higher relative quality of rentals. The positive correlation in Figure 6 implies that the difference in rent-to-price across cities cannot be explained by high rent-to-price cities having higher relative quality of rentals as assumed in Head et al. (2023).

To illustrate this more clearly, Figure 7 plots the estimated density of quality in the rental and owner-occupied sectors by city for 2014. Before discussing the patterns, I briefly explain how these densities are constructed.

Recall that quality in each city is normalized to equal the price rank as of 2014 in the city. Further, $g_c(q)$ is the estimate of the 2014 price rank in city c for a property with 2014 rent rank q. It follows that I can obtain the density of quality in the rental sector in 2014 by evaluating $g_c(q)$ over the rent ranks q and estimating a kernel density on the resulting values for $g_c(q)$. Let $g^r_{c,2014}$ represent the density of quality in the rental sector in city c as of 2014 recovered in this manner.

Figure 7: Comparison of 2014 Quality Distributions by Tenure and City

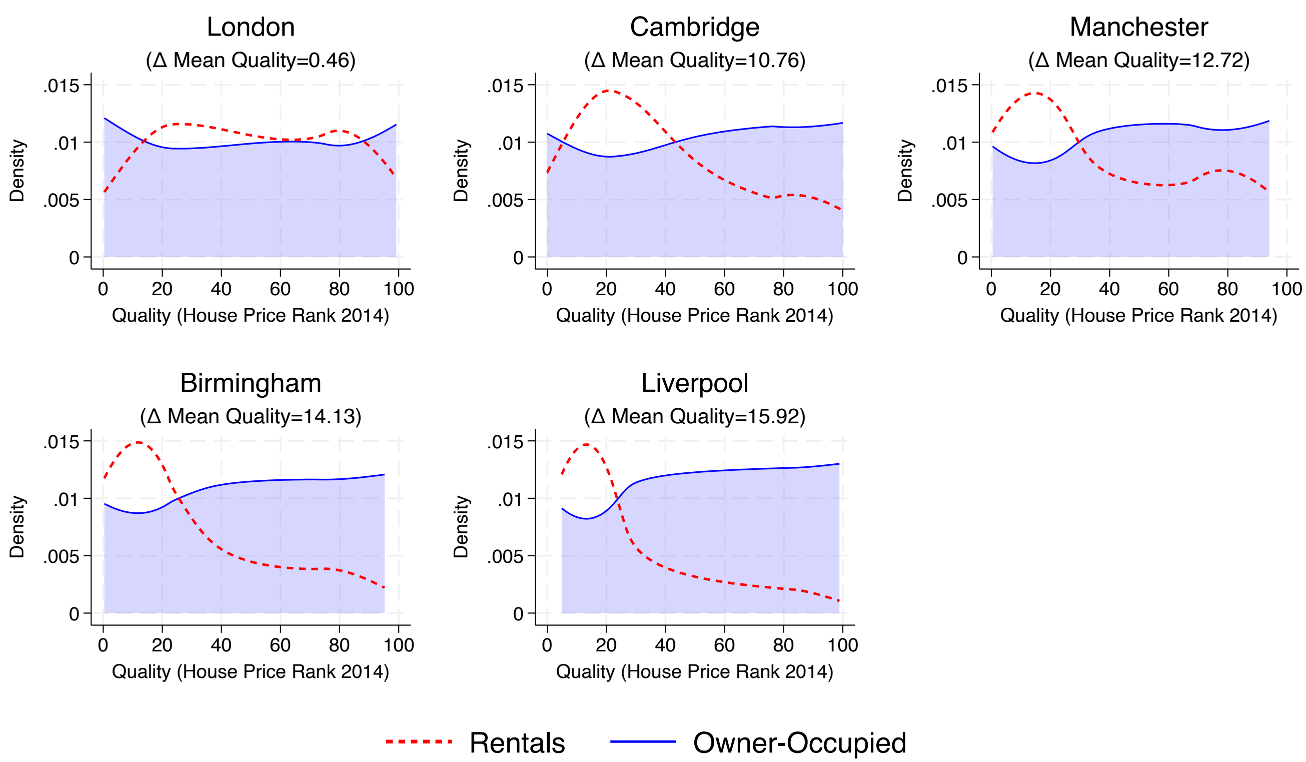

Note: This figure plots the estimated distribution of quality in the rental and owner-occupied sectors by city. The figure also reports the difference in mean quality between the owner-occupied and rental sectors by city.

To back out the implied density of quality in the owner-occupied sector $g^o_{c,2014}$, I rely on the fact that the density of quality of the overall housing stock of a city is a mixture of the densities of quality in the rental and owner-occupied sectors: $g^{tot}_{c,2014}(h) = s_{c,rent} \times g^r_{c,2014}(h) + (1 - s_{c,rent}) \times g^o_{c,2014}(h)$, where quality h is defined as the 2014 price rank; $s_{c,rent} = n_{c,rent}/(n_{c,rent} + n_{c,own})$ is the share of private rentals in c in 2014; and $1 - s_{c,rent}$ is the share of owner-occupied units.[10] This allows me to represent the density of quality in the owner-occupied sector as:

$$g^o_{c,2014}(h) = \frac{g^{tot}_{c,2014}(h) - s_{c,rent} \times g^r_{c,2014}(h)}{1 - s_{c,rent}} \tag{1}$$

The density in Equation 1 is well-defined so long as $g^{tot}_{c,2014}(h) > s_{c,rent} \times g^r_{c,2014}(h)$ and $1 - s_{c,rent} > 0$. Both of these conditions are satisfied. In the empirical implementation, I approximate these continuous densities on a discrete grid. Specifically, I compute the density $g^o_{c,2014}(h)$ over a grid of quality levels ($h = 1, ..., 100$). I require three variables to compute Equation 1. First, I proxy for the 2014 private rental share $s_{c,rent}$ using the analogous share in 2011 obtained using the 2011 Census. Second, for the density in the rental sector, I use estimates of $g^r_{c,2014}$ obtained using the procedure outlined above.[11] Finally, since the quality of the overall housing stock is normalized to equal the 2014 price rank, the density of quality of the housing stock in 2014 is uniform by construction. Since the evaluation is done on a 100-point

[10] Note that social housing is excluded from the analysis throughout the paper.

[11] Given the discretization, I evaluate $g^r_{c,2014}(h)$ at the 100 grid points and rescale the resulting values so that they sum to one, ensuring that the discretized version of the rental density is properly normalized.

grid, the probability mass for each point in the grid is 1/100. As a result, I set $g^{tot}_{c,2014}(h) = 1/100$ in Equation 1.

Figure 7 plots kernel density estimates of $g^{r}_{c,2014}$ in red and $g^{o}_{c,2014}$ in blue. The figure also reports the difference in mean quality between the owner-occupied and rental sectors by city. These comparisons confrim our earlier inference that the difference in rent-to-price ratios across cities cannot be explained by high rent-to-price cities having higher relative quality of rentals. Figure 7 reveals that the opposite is true. The relative quality of rentals (compared to owner-occupied properties) tend to be higher in low rent-to-price cities such as London.

If not differences in relative quality by tenure, what else might explain the systematic variation in the mean rent-to-price ratios across cities? I next show evidence which suggests that differences in the mean rent-to-price ratios across cities arise due to differences in landlords' opportunity costs of capital.

1.3.2 Rent-to-Price Ratio and Capital Gains Covary Negatively Across Cities But Positively Over Quality Segments Within Cities

Capital gains from house price appreciation is an important determinant of the opportunity costs of capital. In Figure 8 I investigate how the rent-to-price ratio varies with annualized capital gains computed using properties that were sold both in 2010 and 2018.

The figure on the left plots the relationship between the average rent-to-price ratio and the average annualized capital gains in different cities. The relationship

Figure 8: Rent-to-Price Ratio and Capital Gains Covary Negatively Across Cities But Positively Over Quality Segments Within Cities

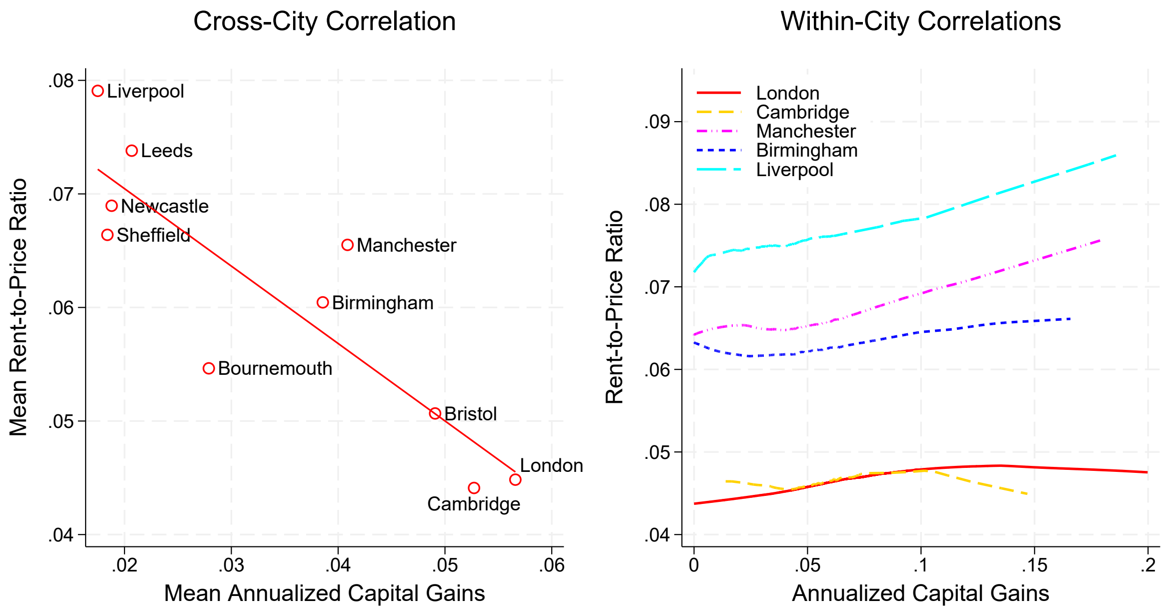

Note: This figure plots the relationship between the rent-to-price ratio and capital gains both across and within cities. The capital gains are computed using data on properties that were sold both in 2010 and 2018. The figure on the left plots the mean rent-to-price in a city against the mean annualized capital gains in that city. The figure on the right presents lowess-smoothed plots for each city of the conditional expected rent-to-price ratio for properties with a given level of annualized capital gains. To construct the figure on the right, I impute the rent-to-price ratio for each property in the repeat sales sample given the 2010 sales price. The imputation is based on an OLS regession of the rent-to-price ratio on a price polynomial, with the degree of the price polynomial selected using five-fold cross-validation.

is negative providing further evidence that the difference in the rent-to-price ratios across cities is driven by differences in the opportunity costs of capital.[12]

The figure on the right plots the relationship between the rent-to-price ratio and annualized capital gains within cities. The plots are based on the sample of properties that were transacted both in 2010 and 2018. I compute the annualized

[12]The negative correlation across cities is consistent with existing studies in the literature. Demers and Eisfeldt (2022) show that capital gains co-vary negatively with net rental yields across US cities. Amaral et al. (2025) find a similar cross-city relationship when looking at 15 OECD countries.

capital gains for each property based on the prices in 2010 and 2018. I also need to impute rent-to-price ratios for the properties based on 2010 prices. I achieve this in two steps.

In the first step, I estimate OLS regressions by city to model the relationship between the rent-to-price ratio and property prices during the period 2014-2018. Using the mappings discussed in Section 1.2.1, I impute prices $\widehat{price}_{qcy}$ and the rent-to-price ratio $\widehat{RTP}_{qcy}$ for each rent percentile rank q of the rental stock by city and year. Using these imputed values, I estimate a flexible polynomial regression of the form $\widehat{RTP}_{qcy} = f_c(\widehat{price}_{qcy}) + \varepsilon_{qcy}$, where f_c is a polynomial function that varies by city. The polynomial degree for each city is chosen by five-fold cross-validation. Appendix Table A.3 reports the optimal polynomial degree for each city selected by cross-validation. The fitted mapping $\widehat{f}_c(p)$ gives the expected rent-to-price ratio given price p. In the second step, I use these estimated mappings to impute the rent-to-price ratio as of 2010 given 2010 prices. This imputation assumes that the relationship between the rent-to-price ratio and price in 2010 was the same as in the period 2014-2018 used to estimate $\widehat{f}_c(p)$.

The lowess-smoothed plots in Figure 8 present the relationship between the imputed rent-to-price ratio and the annualized capital gains for the properties in the repeat-sales sample. The negative correlation documented in the cross-city comparison disappears, with several cities showing evidence of a positive correlation between the rent-to-price ratio and capital gains. The lack of a negative correlation implies that within-city differences in the rent-to-price ratio are more likely to be driven by factors other than capital gains.

The within-city correlations suggest that for properties transacted in 2010 and 2018, capital gains either tend to decrease in prices or are uncorrelated with prices.[13] Appendix Figure A.4 shows that a similar pattern holds when capital gains are computed using properties that were transacted in 2014 and 2018. While reassuring, care must be taken when interpreting these correlations since this pattern may not hold in all periods. For instance, focusing on a sample of transactions spanning 2009-2014, Halket et al. (2020) report that capital gains tend to be higher closer to central London, where dwellings tend to be more expensive. Looking to the US, Landvoigt et al. (2015) find that in San Diego, capital gains were negatively correlated with prices during the boom in 2000-2004, but the relationship reversed during the bust that followed.

Finally, to check whether the patterns are sensitive to the order of the polynomial regressions used to impute the rent-to-price ratio, Appendix Figure A.5 reconstructs the within-city patterns in Figure 8 based on rent-to-price ratios which are imputed using a specification which is linear in price. While the correlations change slightly, the use of this simpler specification does not impact the qualitative conclusions.

Taken together, I interpret the evidence as suggesting that capital gains may play a key role in explaining differences in the rent-to-price across cities, but it may be less important for explaining rent-to-price ratios across quality segments within cities. Motivated by these patterns, in the quantitative model I allow landlords in different cities to face different unobserved user costs. Using counterfactual simulations I show that differences in user costs across cities indeed explain much of the cross-

[13]This follows from the observation that the rent-to-price ratio is decreasing in prices within cities.

city variation in rent-to-price ratios. On the other hand, the within-city variation in the rent-to-price ratio is explained by within-city heterogeneity in landlord wealth endowments which impacts the costs facing landlords in different quality segments.

1.3.3 Depreciation, Risk and the Rent-to-Price Ratio

The patterns documented above suggest that differences in capital gains help explain cross-city variation in rent-to-price ratios. However, two additional mechanisms, differences in effective depreciation rates and differences in risk premia, can also affect the user costs faced by landlords.

First, the higher prices in cities such as London imply that land accounts for a larger share of property values, suggesting lower effective depreciation rates. Second, differences in risk can also matter. Amaral et al. (2025) document that housing assets in large, high-demand cities are perceived as safer and more liquid, implying lower required risk premia for investors. In the standard user-cost formulation (e.g. Poterba, 1992; Poterba and Sinai, 2008), higher depreciation and risk premia both increase the cost of owning housing.

Taken together, these patterns suggest that variation in user costs arising from differences in capital gains, depreciation, and risk premia may explain much of the observed cross-city variation in rent-to-price ratios. Motivated by these observations, in the structural estimation in Chapter 2, I allow the net impact of these factors to vary across cities but do not decompose their individual contributions. A full decomposition would enrich the model and enable analysis of a broader range of counterfactuals. I intend to pursue this in future work.

1.3.4 Landlords are Relatively Wealthier in Cities with Lower Mean Rent-to-Price Ratios

Landlord costs can vary due to differences in the extent to which they rely on borrowing to finance their properties. Borrowing can increase costs of capital through two channels. First, it increases exposure to costs of borrowing which may be greater than the cost of financing with cash. Second, increased levels of borrowing can increase interest rates charged by lenders. Borrowing might be higher for landlords in London due to the higher prices of properties in London.

To test this hypothesis, I use data from 2014-2019 from the Wealth and Assets Survey, which records the total outstanding mortgage and total property value of buy-to-let properties. Using these variables, I construct the loan-to-value ratio on buy-to-let properties for each landlord in the sample. Restricting the sample to landlords in England, I then estimate an OLS regression of the form $LTV_{it} = \alpha + \beta \mathbb{1}(i \in \text{London}) + \varepsilon_{it}$, where $\mathbb{1}(i \in \text{London})$ is an indicator for whether landlord i is in London. The estimate for α is 0.324 (SE=0.011), while the estimate for β is 0.022 (SE=0.029). The point estimates indicate that the loan-to-value for landlords in London is approximately two percentage points higher than the loan-to-value for landlords in the rest of England. However, the difference is not statistically significant. Adding year fixed effects reduces the estimate for β to 0.0179 (SE=.028).

The point estimate of a two percentage point difference in loan-to-value is suprisingly small considering the large differences in prices between London and the rest of England. The small difference could be explained by landlords in London having greater wealth than landlords elsewhere. To test this possibility, Figure 9

Figure 9: Difference Between Landlord and Household Median Assets Decreasing in the Mean Rent to Price

Note: This figure plots the difference in median assets between landlords and households against the mean rent-to-price ratio in the city. The figure also plots a line of best fit.

presents a scatter plot of relative landlord assets in each city (compared to households) against the mean rent-to-price ratio in each city. Consistent with the idea that greater landlord wealth allows landlords to avoid high loan-to-value ratios in London, the figure shows that the relative wealth of landlords (compared to households) is indeed higher in London and Cambridge compared to other cities. More generally, the figure shows that landlords are wealthier relative to households in cities with lower rent-to-price ratios.

These results highlight the importance of accounting for inter-city differences in landlord endowments in a structural model. Greater landlord wealth offsets the effect of higher property prices on borrowing. Failing to account for differences in landlord wealth would lead to overstated loan-to-values in cities like London where both prices and endowments are high. Accordingly, in Chapter 2 I explicitly allow for variation in landlord endowments across cities and allow these endowments to shape rental supply.

1.4 The Need for a Model

The systematic heterogeneity documented in this chapter suggests landlord supply constraints may play an important role in shaping the housing market and in generating the differences in equilibrium outcomes observed both within and across cities. While suggestive, these patterns are not sufficient to pin down the role of landlord supply constraints due to two reasons. First, they are based on correlations between endogenous outcomes including rents, prices, and the sorting of households to housing. Second, these stylized facts do not account for differences in demand both within and across cities arising from heterogeneity in household preferences and endowments of income and assets. To overcome these issues, in Chapter 2 I develop an equilibrium model of the housing market which allows me to quantify the impact of landlord supply constraints while accounting for rich heterogeneity in household preferences, endowments, and constraints.

Chapter 2

An Equilibrium Model of the Markets for Rental and Owner-Occupied Housing*

2.1 Introduction

To understand the role of landlord supply constraints in generating the observed differences in housing markets across cities, I develop a two-sided assignment model of local housing markets which builds on Epple et al. (2020)'s equilibrium model of metropolitan housing markets with vertically differentiated housing. I enrich their model in several important directions. These extensions allow the model to endogenously generate the patterns of quality segmentation, ownership rates, and the rent-to-price functions observed in the data.

On the demand side, heterogeneous households simultaneously choose a continuous level of housing quality and decide whether to rent or own. Their choices are influenced by unobserved preferences for quality and owner-occupation,

*Chapters 1 and 2 use Crown-copyright statistical data from the Office for National Statistics (ONS) andthe Consumer Data Research Centre (project ID CDRC 1015-01, ES/L011840/1; ES/L011891/1.). These chapters also use data available under Open Government Licence v3.0. from HM Land Registry, the Department for Levelling Up, Housing & Communities (English Housing Survey), and the Valuation Office Agency. The use of the data does not imply the endorsement of the data provider in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

observed income and asset endowments, and borrowing constraints that limit the quality range that the household can access if they choose to owner-occupy.

Households sort positively to house quality based on income. However, the correlation between income and quality may not be perfect for several reasons. First, the heterogeneity in unobserved preferences implies that households with the same income may choose different levels of quality. Second, the choice sets facing households may vary independently of income due to the impact of assets on borrowing constraints. Third, differences between the user costs of owner-occupation and rents may drive a wedge in the choice of quality by tenure conditional on income.

I depart from the literature by allowing unobserved user costs of owner-occupation to be different than equilibrium rents. The literature conventionally treats user costs as equivalent to rents (Poterba, 1984, 1992), an equivalence that only holds when households face neither borrowing constraints nor preferences over tenure. When renting and owning are perfectly substitutable, any cost wedge between them is arbitraged away.[14] In my model, this rental equivalence breaks down since both borrowing constraints and tenure preferences are present. Preferences for tenure create a wedge between rents and user costs, and borrowing constraints prevent households from fully offsetting these differences even when they would otherwise prefer to own.

On the rental supply side, atomistic landlords choose the profit-maximizing house quality to own and let out in the local rental market. The choice depends on their observed asset endowments and unobserved opportunity costs of capital. The

[14]When households face neither borrowing constraints nor tenure preferences, rental equivalence must hold; otherwise, all households choose the cheaper option, and only one sector survives.

unobserved opportunity cost of capital captures costs of maintenance, depreciation, borrowing, taxes and a risk premium net of capital gains (Poterba, 1984, 1992). The evidence shows that the cost of borrowing is convex in the loan-to-value ratio in the UK (Benetton, 2021). I therefore define landlords' costs of capital to be a convex function of the share of property financed by borrowing. The convexity of costs with respect to borrowing implies that marginal profits with respect to quality are greater for wealthier landlords.[15] This leads to a positive assignment of landlords to housing quality based on wealth, as wealthier landlords stand to gain more from operating in higher quality segments. Hence, the supply of rentals in each quality segment crucially depends on the distribution of landlord assets and the shape of their cost function.

Equilibrium in the model is characterized by two market clearing conditions for each quality segment. First, the demand for rental services must equal the supply by landlords in each quality segment. Second, the demand for ownership by landlords and households must equal the available stock of housing in each quality segment. Prices and rents are determined by the marginal household and landlord within each quality segment. The schedule of prices and rents are non-linear over quality because marginal households and landlords vary across quality segments (Landvoigt et al., 2015; Epple et al., 2020).

[15] I assume that landlords only borrow when they lack sufficient wealth to purchase the property outright, consistent with UK survey evidence showing landlords prefer to avoid borrowing when possible (Scanlon and Whitehead, 2016). Consequently, wealthier landlords borrow less than their less wealthy counterparts across all quality levels. Because costs are convex in borrowing, wealthier landlords tend to operate in the less convex part of the cost function, leading to higher marginal profits from operating in higher quality segments.

Quality segmentation arises in the model when the equilibrium distribution of house quality varies across the rental and owner-occupied sectors. The model provides several channels, on both the demand and supply side, through which quality segmentation can be generated. On the demand side, segmentation arises when unobserved preferences for owner-occupation are positively correlated with either preferences for housing quality or household income, both of which increase households' demand for quality. On the supply side, landlords can shape quality segmentation by influencing the wedge between rents and households' user costs of owner-occupation. The wedge in costs affects segmentation because households' optimal choice of quality depends on rents when renting and on user costs when owning.

I estimate the structural parameters of the model using a method of moments estimator exploiting data for nine markets: London, Manchester, and Birmingham for the years 2014-2016. The estimation takes as inputs the joint distribution of wealth and income for households as well as the marginal distribution of wealth for landlords, which I recover from the UK Wealth and Assets Survey. The estimation imposes market clearing for each quality segment in each market. The key moments include for each market the marginal distributions of prices and rents, the rent-to-price ratios, ownership rates, quality segmentation, and correlations between household endowments and rents and prices, which I recover using data from Zoopla, the UK Wealth and Assets Survey and the UK Household Longitudinal Survey.

I rely on two sources of identification to separately identify the household's preferences from the landlord's cost parameters. First, I exploit variation in equilibrium outcomes and the distributions of household and landlord endowments across cities and years. The model maps household and landlord endowment distributions to equilibrium outcomes, with household parameters moderating the effect of household endowments on the equilibrium and landlord parameters shaping the effect of landlord endowments. Independent variation in household and landlord endowment distributions across cities and years thus helps identify household and landlord parameters.[16] Second, the estimation uses several demand-side moments that do not depend on landlord parameters conditional on prices and rents. These include ownership rates by age, correlations between house quality and household endowment by age, and quality segmentation by tenure. These exclusion restrictions further help me disentangle household preferences from landlord costs.

The estimated model successfully replicates the empirical rent and price distributions as well as the patterns of quality segmentation observed in the data. Estimates of the household demand parameters show that heterogeneity in unobserved preferences is key to explaining differences in the choice of quality and the decision to rent or own by age, even after conditioning for differences in income and assets. However, unobserved household preferences play a smaller role in explaining differences in sorting patterns across cities. Estimates for the landlord's cost function exhibit large differences across cities. This suggests that landlord supply constraints mostly account for the observed differences in quality segmentation and rent-to-price ratios across cities.

[16] See Heckman et al. (2010) on identification of hedonic models using data from multiple markets.

I conduct three counterfactual experiments that illustrate that heterogeneity in landlord supply constraints is key to explaining differences in housing markets both across cities and across quality segments within cities. In the first counterfactual, I show that within-city heterogeneity in landlord wealth almost completely explains the negative slope of the rent-to-price function with respect to prices within cities—a robust feature of housing markets that has been documented both in the UK and the US.[17] This suggests that within-city wealth disparities among landlords are critical in shaping rental affordability over the distribution of quality. I find that reducing the variance of the landlords' wealth distribution by 90% almost completely flattens the equilibrium rent-to-price function. This occurs because a reduction in the variance of the wealth distribution increases the relative wealth of landlords operating in lower quality segments which in turn reduces their costs relative to landlords in higher quality segments.[18] Landlords respond to the relative reduction in costs by reducing rents in lower quality segments. This flattens the rent-to-price function.

In the second counterfactual, I show that the cost advantage enjoyed by landlords in London over those in Birmingham accounts for much of the cross-city difference in the housing market equilibrium. To quantify this effect, I set the net user costs for landlords in London to Birmingham levels and solve for the resulting equilibrium in London. This reduces the gap in quality segmentation between London and Birmingham by over 60%. The change in quality segmentation

[17] Studies which document the negative slope of the rent-to-price include Verbrugge (2008); Verbrugge and Poole (2010); Bracke (2015); Halket et al. (2020).

[18] Recall that in the model landlords sort positively to house quality based on assets. This implies that a reduction in variance increases wealth in lower segments and reduces it in higher segments. This tends to reduce relative costs in lower segments since costs are increasing in borrowing which is in turn falling in wealth.

is driven by a 42% reduction in the rent-to-price gap between these two cities and a 67% reduction in the user cost-to-rent gap. These results indicate that quality segmentation varies across cities largely due to differences in rental supply. In cities such as London landlords are willing to accept lower rents from tenants and pay higher prices to acquire properties which reduces rents both in absolute terms and relative to the user costs of owner-occupation. The reduction in the costs of renting allows a greater share of renters to access housing in higher-quality segments, thereby reducing quality segmentation.

To explore the welfare implications of the heterogeneity in landlord supply constraints, I analyze the compensating variation associated with increasing landlord costs in London to Birmingham levels. I compute the compensating variation numerically for the distribution of households in London. I find that the cost advantage enjoyed by landlords in London increases the welfare of the median renter in London by £1,448 per year and lowers the welfare of the median owner-occupier by £2,139 per year. These welfare effects reflect the trade-off arising from the reduction of rents relative to prices as well as the increase in the relative quality of rentals. Additionally, I find that the magnitude of the welfare effects increase in household income. This occurs because in the model a reduction in landlord costs increases the profit-maximizing quality for marginal landlords in each quality segment. This causes the assignment of landlord assets to houses to cascade upwards magnifying the supply effect in higher quality segments.

In the third counterfactual, I impose a 30% tax on rental income in each city. I find that the tax elasticity of the equilibrium rent-to-price function is up to five times greater in Birmingham than in London. This shows that the same tax can lead to vastly different outcomes depending on the composition of households and landlords in the local housing market. The lower pass-through of the tax in London reflects the greater cost advantages and higher profits enjoyed by landlords in London, enabling them to absorb more of the tax burden compared to their counterparts in Birmingham. These results suggest that a progressive tax structure—one that increases with landlord income—may distort the housing market less than a flat tax, as it would better account for the varying capacities of landlords across cities to absorb taxes.

2.1.1 Related Literature

Chapters 1 and 2 of this dissertation are related to a large and growing literature on housing choice. A large strand of the literature models demand for differentiated housing and/or neighbourhoods while abstracting from the choice of whether to rent or own (Bayer et al., 2007, 2016; Galiani et al., 2015; Landvoigt et al., 2015; Diamond et al., 2019; Almagro and Dominguez-Iino., 2020; Calder-Wang, 2020; Epple et al., 2020). Another strand composed mainly of macro-housing papers focuses on the choice to rent or own but assumes that housing is not differentiated (Chambers et al., 2009a,b; Sommer et al., 2013; Binner and Day, 2015; Halket and Pignatti Morano di Custoza, 2015; Sommer and Sullivan, 2018; Kaplan et al., 2020; Han et al., 2021; Greenwald and Guren, 2024). This paper belongs to a nascent third strand which studies housing markets with differentiated housing and both a rental and an owner-occupied sector (Head et al., 2023; Higgins, 2023; Kvaerner et al., 2023).

To my knowledge, Head et al. (2023) is the only other paper that develops a model to examine cross-city differences in the composition of properties between the rental and owner-occupied sectors. While they use the price-to-rent ratio as a proxy for quality differences, I develop a measure of relative quality that is distinct from this ratio, allowing me to explore the link between relative quality and the price-to-rent ratio. In Section 1.2, I show that the price-to-rent ratio does not consistently capture quality segmentation, as the two measures are strongly and negatively correlated across English cities. I develop a model which explains the negative correlation observed in the data.

Several studies have found that the rent-to-price function decreases in price both in the UK and in the US (Verbrugge, 2008; Verbrugge and Poole, 2010; Heston and Nakamura, 2011; Bracke, 2015; Halket et al., 2020). The literature has considered a number of explanations for this phenomenon. Halket and Pignatti Morano di Custoza (2015) show that rental vacancy rates can explain some of the relationship. Halket et al. (2020) argue that this pattern can arise when high price rental properties have worse unobserved characteristics than high price owner-occupied properties. Bracke (2015) shows that the rent-to-price is falling in rent appreciation. However, using data on transaction prices I find evidence suggesting that the rent-to-price function is increasing in capital gains within cities. This paper offers an alternative rental supply-side explanation which almost completely explains the negative slope of the rent-to-price function within cities. I show using counterfactual experiments that the slope of the rent-to-price is linked to the variance of landlord wealth distribution. I find that reducing the variance of landlord wealth by 90% almost completely flattens the rent-to-price function.

This paper also contributes to the measurement of unobserved user costs of owner-occupation. The housing choice literature often assumes that observed rents equal unobserved user costs of owner-occupation. Theoretically the equivalence of rents and user costs holds when households are indifferent between renting and owning and have the ability to arbitrage away cost differences between tenure types (Poterba, 1992, 1984). However, both of these conditions are violated in this paper since the model allows households to prefer owning to renting and incorporates credit constraints which restrict arbitrage. This paper provides a structural approach to estimate unobserved user costs when rental equivalence cannot be assumed. The approach is related to Díaz and Luengo-Prado (2008) and Blow and Nesheim (2009) who also estimate user costs as the shadow cost of owner-occupation. More broadly, the approach complements the literature measuring user costs of homeownership (Bishop et al., 2023).

Finally, I contribute to an emerging literature which investigates heterogeneity in returns to housing. Amaral et al. (2025) show that total returns are lower in larger cities and attribute this to spatial variation in risk. In contrast, the structural model in this paper incorporates a baseline variable cost term for landlords that increases with risk and depreciation and decreases with capital gains. I show that this baseline cost is higher in cities with higher mean rent-to-price ratios. Damen et al. (2025) find that lower-rent properties within cities earn higher returns, and argue that this reflects persistent segmentation in the investor base whereby smaller landlords operate in low-end segments while institutional investors avoid them due to reputational concerns, limiting arbitrage. The assignment model in my paper generates a similar asset-based sorting of investors into quality segments due to financing costs which are convex in the loan-to-value ratio. I show that this sorting induces heterogeneity in landlords costs across quality segments which in turn explains variation in the rent-to-price ratio across these quality segments.

2.2 The Model

2.2.1 Setup

Each market is defined by a unique combination of city c and year t.

Housing—I model housing as a vertically differentiated product. Let h represent the unidimensional index of housing quality. There are $N_{ct}(h)$ houses of quality h in the market. Any house can be rented or owned at the equilibrium rents and prices. Prices $p_{ct}(h)$ evolve non-linearly over house quality. Homeownership credit constraints imply that rents may be different from the user costs of owner occupation. Consequently, I allow rents $r_{ct}(h)$ and user-costs of owner-occupation $c_{ct}(h)$ to have separate non-linear schedules over house quality.

Households—Each market has a distribution of households that vary along four dimensions: income y, assets a, observed characteristics x, and unobserved type i. The share of households with characteristics x equals $s_{ct}(x)$. Households with observed characteristics x draw endowments from the conditional distribution

$F_{ct}(a, y \mid x)$ and belong to unobserved type i with conditional probability $p_{i|x}$. I assume that the distribution of endowments does not depend on the unobserved type after conditioning on observed characteristics i.e., $F_{ct}(a, y \mid x) = F_{ct}(a, y \mid i, x)$ for all i. The distribution of endowments for unobserved type i is then given by

$$F_{ict}(a, y) = \sum_x \left(\frac{s_{ct}(x) p_{i|x}}{\sum_x s_{ct}(x) p_{i|x}} \right) F_{ct}(a, y \mid x) \tag{2}$$

The household's choice is jointly determined by its endowment and unobserved type. The endowment shapes the households' choice set through the budget constraint and homeownership credit constraint. The unobserved type determines the household's preferences for housing quality and homeownership. Conditional on the endowment and unobserved type, households' choices do not depend on observed characteristics x.

Landlords—In each city, landlords own houses and rent out housing services to those households who choose to rent. Landlords vary in their asset endowments. They seek to maximize profits from rents and capital gains net of operating costs. The landlord's endowment shapes the costs they face to operate rentals of varying quality. This in turn impacts the quality of house they choose to own and rent out. Aggregating over the choices of all landlords gives rise to the supply of rental services over the distribution of house quality.

2.2.2 Housing Demand

Each household chooses housing and non-housing consumption z to maximize utility. Housing consumption is defined by the choice of house quality h and the decision to rent or own τ, henceforth referred to as tenure choice. Households of unobserved type i with endowments (a, y) solve the following problem:

$$
\begin{aligned}
&\max_{h,\tau,z} \quad u(h, z \mid \theta_i^h) + \beta_i^\tau + \varepsilon_{i\tau} && \text{subject to} \\
&\quad y = r_{ct}(h) + z && \text{if } \tau = \text{rent} \\
&\left.\begin{aligned} &y = c_{ct}(h) + z \\ &p_{ct}(h) \leq \overline{p}_{ct}(a, y) \equiv \min\left\{a/(1-\phi^{ltv}),\ \phi^{lti} y + a\right\} \end{aligned}\right\} && \text{if } \tau = \text{own}
\end{aligned}
\tag{3}
$$

where θ_i^h is the preference for housing quality for unobserved type i; β_i^τ is the deterministic component of households' preference for tenure τ; $\varepsilon_{i\tau}$ is an idiosyncratic preference for tenure; ϕ^{ltv} and ϕ^{lti} are the maximum loan-to-value and loan-to-income ratios which restrict the size of mortgage that a household can obtain; and $\overline{p}_{ct}(a, y)$ is the maximum house price the household can afford given its endowment.

The choice of tenure determines whether the household pays rents $r_{ct}(h)$ or user costs of owner occupation $c_{ct}(h)$ to reside in a house. The user cost $c_{ct}(h)$ is meant to capture ongoing monetary and opportunity costs associated with owneroccupying a house net of capital gains.[19] The model is flexible in its treatment of rents and user costs and admits a variety of relationships between the two. For instance, the often invoked assumption of rental equivalence arises as a special case when $r_{ct}(h) = c_{ct}(h)$. I depart from the convention of assuming rental equivalence because it is only guaranteed to hold in frictionless settings in which households are indifferent between renting and owning houses of the same quality. Given the central role of credit constraints in this paper, I instead allow rents to vary from user costs of owner-occupation. Differences in rents and user costs allows the model to generate rich trade-offs between ownership and house quality. However, since user costs are unobserved, this added flexibility comes at the cost of having to estimate user costs. I exploit the equilibrium implications of the structural model to recover the unobserved costs as the shadow costs of owner-occupation.

Optimal Choice of Quality by Tenure— The household's choice of quality if it were to rent is given by the level $h^r_{ict}(y)$ which solves the first order condition

$$\frac{\frac{\partial}{\partial h} u(h, y - r_{ct}(h) \mid \theta^h_i)}{\frac{\partial}{\partial z} u(h, y - r_{ct}(h) \mid \theta^h_i)} = r'_{ct}(h) \tag{4}$$

When households do not face binding credit constraints, their choice of quality if they were to own is given by the level $h^{o*}_{ict}(y)$ which solves the first order condition

$$\frac{\frac{\partial}{\partial h} u(h, y - c_{ct}(h) \mid \theta^h_i)}{\frac{\partial}{\partial z} u(h, y - c_{ct}(h) \mid \theta^h_i)} = c'_{ct}(h) \tag{5}$$

[19]These user costs are increasing in costs assocated with taxes, depreciation, maintenance, insurance and opportunity costs of capital as in Poterba (1992, 1984). Expected capital gains reduce user costs.

When rents equal user costs, the two first order conditions coincide. Therefore when costs do not vary by tenure, preferences for quality also do not vary by tenure. On the other hand, differences in rents and user costs generates differences in the optimal levels of quality by tenure. Consider the case when rents are more expensive than the user costs of owning. When this occurs, the marginal rate of substitution for house quality h in Equation 4 is lower than the marginal rate of substitution for h in Equation 5. This implies that the preferred quality will be higher when the household owns, $h^{o^*}_{ict}(y) > h^{r}_{ict}(y)$.

When households do not face binding homeownership credit constraints, their choice of housing exactly reflect their preferences for quality. Therefore, any difference in the preference for quality by tenure is reflected one-for-one in the choice of quality by tenure. However, this does not apply when households are credit constrained. A household is defined as being credit-constrained if the quality it prefers to owner-occupy is priced above the maximum price it can afford given its endowment, $p_{ct}(h^{o^*}_{ict}(y)) > \overline{p}_{ct}(a,y)$. Let $p^{-1}_{ct}\left(\overline{p}_{ct}(a,y)\right)$ represent the highest level of quality that the household can afford to owner-occupy. Given this definition, we can define the quality that a household would actually choose if it decided to owner-occupy as:

$$h^{o}_{ict}(a,y) = \begin{cases} h^{o^*}_{ict}(y) & \text{if } p_{ct}(h^{o^*}_{ict}(y)) \leq \overline{p}_{ct}(a,y) \\ p^{-1}_{ct}\left(\overline{p}_{ct}(a,y)\right) & \text{otherwise} \end{cases} \tag{6}$$

Tenure Choice—Households choose the tenure which maximizes their utility. Let $u^r_{ict}(y)$ and $u^o_{ict}(a,y)$ be the household's utility from renting and owning their constrained optimal quality respectively:

$$u^r_{ict}(y) \equiv u\big(\, h^r_{ict}(y), \qquad y - r_{ct}(h^r_{ict}(y)) \quad \mid \theta^h_i \big) \tag{7}$$

$$u^o_{ict}(a,y) \equiv u\big(\, h^o_{ict}(a,y), \quad y - c_{ct}(h^o_{ict}(a,y)) \mid \theta^h_i \big) \tag{8}$$

A household chooses to owner-occupy $h^o_{ict}(a,y)$ if it derives greater utility from doing so:

$$u^o_{ict}(a,y) + \beta^{own}_i + \varepsilon_{iown} \geq u^r_{ict}(y) + \varepsilon_{irent} \tag{9}$$

It rents quality $h^r_{ict}(a,y)$ otherwise. The decision to rent or own depends on a quality-ownership trade-off encapsulated in Equations 6-9. The trade-off arises when households are sufficiently credit constrained. Recall that $p_{ct}(h^r_{ict}(y))$ is the price of the house that the household would rent. When $p_{ct}(h^r_{ict}(y)) \leq \overline{p}_{ct}(a,y)$ there is no trade-off because the household can afford to own a house that is at least as good as the house it would rent. The household need not forgo quality to become an owner-occupier. However, a trade-off arises when the maximum price the household can pay falls below the price of the house it would rent, $p_{ct}(h^r_{ict}(y)) > \overline{p}_{ct}(a,y)$. When this occurs, the household cannot afford to own the house it would rent. It must forgo $h^r_{ict}(y) - h^o_{ict}(a,y)$ units of quality to become an owner-occupier. Therefore, when households are sufficiently credit constrained, the choice of tenure depends on the extent to which they value ownership over quality.

The household's preference for quality depends on the utility function $u(h, z \mid \theta_i^h)$. I defer a discussion of the specification of u to Section 2.3. The value from owning a property depends on cost differences by tenure and tastes for ownership. Cost differences influence the value of ownership through the budget constraint. When renting is more expensive, households that do not face binding credit constraints can increase both their housing and non-housing consumption by owning instead of renting. This increases the value of owning. The taste for tenure impacts the value of owning by shifting the household's indifference curves associated with owning relative to renting. The taste for tenure has a deterministic component β_i^τ and an idiosyncratic component $\varepsilon_{i\tau}$. These tastes capture preferences for tenure that arise due to factors not related to user cost, prices, or house quality. For instance, some households may prefer to have the option to renovate or modify their dwelling. Others may prefer the stability that comes from not having to renew a lease every year. The taste terms β_i^τ and $\varepsilon_{i\tau}$ capture preferences for tenure which arise due to such considerations.

I assume that the idiosyncratic preferences for tenure $\varepsilon_{i\tau}$ are independently and identically distributed Type I Extreme Value with variance σ_{logit}. Moreover I normalize β_i^{rent} to zero. Under this normalization β_i^{own} captures differences between households' non-idiosyncratic tastes for owning and renting, $\beta_i^{own} - \beta_i^{rent}$. Given the assumption on ε, the share of type i households with endowments (a, y) who owner-occupy houses of quality $h_{ict}^o(a, y)$ equals $o_{ict}(a, y)$:

$$o_{ict}(a, y) = \frac{\exp\{(u_{ict}^o(a, y) + \beta_i^{own})/\sigma_{logit}\}}{\exp\{u_{ict}^r(y)/\sigma_{logit}\} + \exp\{(u_{ict}^o(a, y) + \beta_i^{own})/\sigma_{logit}\}} \tag{10}$$

The remaining share $1 - o_{ict}(a, y)$ choose to rent $h^r_{ict}(y)$. The variance parameter σ_{logit} determines the sensitivity of tenure choice to differences in the non-idiosyncratic component of utility across tenure. As σ_{logit} converges to zero, the tenure choice becomes increasingly deterministic and the tenure share converges to the indicator function $\mathbb{1}\left(u^o_{ict}(a, y) + \beta^{own}_i \geq u^r_{ict}(y)\right)$. Conversely, as σ_{logit} converges to ∞, the choice of tenure becomes increasingly random and the tenure share converges to $1/2$.

Demand for Quality—We can obtain the aggregate demand for quality by owner-occupiers and renters of unobserved type i by integrating over the joint distribution of income and assets for the unobserved type:

$$D^o_{ict}(h, p_{ct}, r_{ct}) = \int_{a,y} \quad o_{ict}(a, y) \quad \times \mathbb{1}\left(h^o_{ict}(a, y) \leq h\right) \; dF_{ict}(a, y) \tag{11}$$

$$D^r_{ict}(h, p_{ct}, r_{ct}) = \int_{a,y} \left(1 - o_{ict}(a, y)\right) \times \mathbb{1}\left(h^r_{ict}(y) \leq h\right) \quad dF_{ict}(a, y) \tag{12}$$

To get the demand by tenure for households with observed characteristic x, we can weight Equations 11 and 12 by the conditional distribution of types given the observed characteristic and sum over all types:

$$D^o_{ct}(h, p_{ct}, r_{ct} \mid x) = \sum_i D^o_{ict}(h, p_{ct}, r_{ct}) \times \pi_{i|x} \tag{13}$$

$$D^r_{ct}(h, p_{ct}, r_{ct} \mid x) = \sum_i D^r_{ict}(h, p_{ct}, r_{ct}) \times \pi_{i|x} \tag{14}$$

Finally, weighting Equations 13 and 14 by the share of households with characteristics x and summing over all x gives the aggregate demand for quality h by tenure.

$$D^o_{ct}(h, p_{ct}, r_{ct}) = \sum_x D^o_{ct}(h, p_{ct}, r_{ct} \mid x) \times s_{ct}(x) \tag{15}$$

$$D^r_{ct}(h, p_{ct}, r_{ct}) = \sum_x D^r_{ct}(h, p_{ct}, r_{ct} \mid x) \times s_{ct}(x) \tag{16}$$

2.2.3 Landlord's Problem

Each market has a distribution of landlords with endowments drawn from $G_{ct}(a)$. Landlords are otherwise identical. Each landlord owns and rents out a single house. They choose the optimal house quality based on the conditional expected profits from buying and renting out a house in the current period and selling it the following year given their endowment. The expected profit depends on rents, costs, and expected price appreciation, $\pi_{ct}(h) = \mathbb{E}[p_{ct+1}(h)/p_{ct}(h)]$.

Landlords finance their house purchase using a mix of assets and loans. The financing mix determines the loan-to-value ratio $\phi^a_{ct}(p_{cth})$, which depends on endowments a and endogenous prices $p_{ct}(h)$ for quality h.[20] I assume that $\phi^a_{ct}(p_{cth})$ depends on h only through the impact of h on price; conditional on p_{cth}, the loan-to-value ratio does not vary with housing quality h.

The landlord's costs is variable and proportional to the current price of the property. Let $\mu_{ct}(h, a)$ denote the variable cost for quality h for a landlord with

[20] I use $p_{ct}(h)$ and p_{cth} interchangeably throughout the text.

endowment a:

$$\mu_{ct}(h,a) = 1 + \psi_{ct}(h) + \left(1 - \phi^a_{ct}(p_{cth})\right) \times i^o_{ct} + \phi^a_{ct}(p_{cth}) \times i^b_{ct}(\phi^a_{ct}(p_{cth}))$$

The term $1 + \psi_{ct}(h)$ captures the components of variable cost that are not related to the choice of financing. It is increasing in physical depreciation, maintenance, taxes, and risk premia and decreasing in expected capital gains.[21,22] An increase in the loan-to-value ratio ϕ increases the landlord's exposure to the opportunity cost of borrowing $i^b_{ct}(\phi)$ and reduces exposure to the opportunity cost of cash i^o_{ct}. I assume that $i^b_{ct}(\phi) > i^o_{ct}$ and that $i^b_{ct}(\phi)$ is increasing and convex in the loan-to-value ratio. Landlords therefore prefer financing their purchases with cash and this preference becomes more pronounced as the loan-to-value ratio increases. The landlord only borrows when the price of the house is greater than the value of their assets. This implies that the loan-to-value ratio equals $\phi^a_{ct}(p_{cth}) = 1 - a/p_{ct}(h)$ for landlords who need to borrow. Since the cost of borrowing $i^b_{ct}(\phi)$ is convex in the loan-to-value ratio ϕ, the overall variable cost is also convex in ϕ.

[21]Conceptually, the non-financing component can be written explicitly as $1 + \psi_{ct}(h) = 1 - \pi_{ct}(h) + \delta_{ct} + \beta_{ct} + v_{ct}$, where $\pi_{ct}(h)$ denotes expected capital gains; δ_{ct} is the rate of physical depreciation; β_{ct} is the risk premium associated with operating as a landlord; and v_{ct} is defined broadly to capture taxes and all other costs borne by the landlord which scale with prices. In the UK context, the taxes in v_{ct} can represent the Stamp Duty Land Tax. v_{ct} may also include ground rents for leasehold properties and insurance costs, however, it excludes capital gains taxes which depend on expected price appreciation. Capital gains taxes can be introduced by replacing the $(1 - \pi_{ct}(h))$ term with $(1 - \pi_{ct}(h)) \times (1 - tax^{CG}_{ct})$.

[22]As discussed in Section 2.3.2, in the estimation I use a reduced-form specification to model landlord variable costs. In this specification, the non-financing component $1 + \psi_{ct}(h)$ is absorbed by an intercept which is allowed to vary by city. Hence, the estimation is agnostic about how much each of the various non-financing components (expected capital gains, depreciation, taxes, risk premia, etc.) contribute to variable costs.

The specification of $\phi^a_{ct}(p_{cth})$ implies that, for a given a level of assets a, an increase in price $p_{ct}(h)$ raises the loan-to-value ratio associated with owning and operating a property of quality h. A limitation of this specification is that it does not allow landlords to choose how much they borrow conditional on their assets and current prices.

I assume that landlords fully finance their property purchases with cash whenever possible and only resort to borrowing when their assets cannot fully cover the purchase. This assumption aligns with recent UK survey evidence showing that, among debt-free landlords, the most common reason by far for being mortgage-free is a preference to purchase outright (Scanlon and Whitehead, 2016). This assumption will not hold exactly in practice, since landlords may seek to diversify and hold other assets. As a result, the model tends to overstate the share of assets allocated to housing and understate mortgage borrowing and related costs. The model may therefore understate the sensitivity of landlords' costs to changes in financing conditions. A natural next step would be to endogenize mortgage leverage by allowing landlords to decide how much to borrow and how to allocate their wealth across housing and alternative assets. I intend to pursue this in future work.

Landlord's Choice—The expected profit from houses of quality h for landlords with endowments a is given by

$$\Pi_{ct}(h \mid a) = r_{ct}(h) - \mu_{ct}(h, a) p_{ct}(h) \tag{17}$$

The first order condition for the landlord's optimal choice of quality is

$$\frac{\partial r_{ct}(h)}{\partial h} = \frac{\partial p_{ct}(h)}{\partial h}\mu_{ct}(h,a) + \frac{\partial \mu_{ct}(h,a)}{\partial h}p_{ct}(h) \tag{18}$$

Let $s_{ct}(a, r_{ct}, p_{ct})$ be the house quality which satisfies the first order condition for landlords with endowment a. Then the aggregate supply of rentals of quality h is given by

$$S_{ct}(h \mid r_{ct}, p_{ct}) = \int_a \mathbb{1}\{s_{ct}(a, r_{ct}, p_{ct}) \leq h\}\, dG_{ct}(a) \tag{19}$$

To analyze whether landlords sort systematically to different segments of the housing distribution, I next consider the cross-derivative of the landlord's profit function:

$$\frac{\partial^2 \Pi_{ct}(h \mid a)}{\partial h \partial a} = -\left(\frac{\partial p_{ct}(h)}{\partial h}\frac{\partial \mu_{ct}(h,a)}{\partial a} + \frac{\partial^2 \mu_{ct}(h,a)}{\partial h \partial a}p_{ct}(h)\right) > 0 \tag{20}$$

I verify in Appendix Section B.1 that the cross-derivative of the profit function in Equation 20 is positive. The positive cross derivative implies that the marginal profit from an increase in quality is increasing in a. When marginal profits are increasing in a, the indirect profit function satisfies a single-crossing property in the quality-rent space. Therefore there is positive sorting between landlords and house quality based on endowments: i.e., the landlord's choice of housing quality $s_{ct}(a, r_{ct}, p_{ct})$ is increasing in their endowment a. This allows us to define the inverse function $a_{ct}(h) = s_{ct}^{-1}(a, r_{ct}, p_{ct})$ which gives the endowment of the landlord who rents out a house of quality h. Evaluating the landlord's first order condition at the assigned quality-endowment pair $(h, a_{ct}(h))$ gives the gradient of the rent function with respect to quality at h:

$$\frac{\partial r_{ct}(h)}{\partial h} = \frac{\partial p_{ct}(h)}{\partial h}\mu_{ct}(h, a_{ct}(h)) + \frac{\partial \mu_{ct}(h, a_{ct}(h))}{\partial h}p_{ct}(h) \tag{21}$$

With the rent gradient in hand, we can compute the level of rents for a given quality level by integrating over the distribution of house quality

$$\begin{aligned} r_{ct}(h_j) &= r_{ct}(0) + \int_{h=0}^{h_j} \frac{\partial r_{ct}(h)}{\partial h} dh \\ &= r_{ct}(0) + \int_{h=0}^{h_j} \frac{\partial p_{ct}(h)}{\partial h}\mu_{ct}(h, a_{ct}(h)) + \frac{\partial \mu_{ct}(h, a_{ct}(h))}{\partial h}p_{ct}(h) \; dh \end{aligned} \tag{22}$$

Equation 22 shows that the level of rents depends on the rent derivatives and a constant $r_{ct}(0)$ which represents the baseline level of rents in the market.

2.2.4 Housing Supply

We have previously derived household demand for housing services by tenure and quality, as well as landlord supply of rental services by quality. To complete the model, it is essential to specify the supply of housing quality across different markets. Let m and m' represent two distinct markets. I express the relative housing supply by quality between these two markets using the following recursive relationship:

$$N_{m'}(h) = f^n \left(N_m(h), \kappa_{m'}, \kappa_m \mid \zeta\right) \tag{23}$$

Here, the relative stock of housing of quality h in markets m and m' is influenced by the relative population sizes, $\kappa_{m'}$ and κ_m, in the two markets. The parameter ζ captures factors that moderate how the housing stock responds to population differences across markets.

Suppose $N_0(h)$ represents the housing stock in a baseline market. The recursive specification allows us to compute the housing stock in all other markets, provided $N_0(h)$ is known. This property proves useful in the estimation process, as $N_0(h)$ can be directly measured in the data given a normalization in the base market. This enables me to recover the quality distribution in all other markets given this normalization. I elaborate on this approach in Section 2.3.

2.2.5 Equilibrium

Equilibrium in market ct is characterized by the vector of rents and prices (r_{ct}, p_{ct}) which satisfy the following market clearing conditions for each quality level h:

$$N_{ct}(h \mid N_{ct-1}, p_{ct}, p_{ct-1}) = D^o_{ct}(h, p_{ct}, r_{ct}) + D^r_{ct}(h, p_{ct}, r_{ct}) \tag{24}$$

$$D^r_{ct}(h, p_{ct}, r_{ct}) = S_{ct}(h, p_{ct}, r_{ct}) \tag{25}$$

Equation 24 states that the overall demand for housing of quality h by renters and owner-occupiers must equal the stock of such housing. Equation 25 states that the demand for quality h by renters must equal the supply of rental services in the quality segment by landlords.

2.3 Estimation and Identification

2.3.1 Discretizing the Quality Space

The estimation involves solving the equilibrium conditions in Equation 24 and 25 numerically. Since prices and rents are non-linear over quality, I need to numerically solve for the value of prices and rents at each quality level h. I discretize the quality space into a grid of J points $h_1, ..., h_J$ to make the estimation feasible. The discretization requires that we work with the discrete analogs of the choices and quantities defined in Section 2.2. I discuss the discrete analogs below.

I start with the landlord's case because it is more straightforward. Optimal behaviour of the landlord implies that there exist J asset thresholds $a_{1ct} < ... < a_{Jct}$ such that all rental properties of quality h_j are operated by landlords with assets in the range $a_{jct} \leq a < a_{j+1ct}$. The asset threshold a_{jct} is implicitly defined by the asset value which makes the landlord indifferent between qualities h_j and h_{j+1}: $\Pi_{ct}(h_j \mid a) = \Pi_{ct}(h_{j+1} \mid a)$. Given the thresholds, the supply of rental properties in quality segment h_j is given by $G_{ct}(a_{j+1ct}) - G_{ct}(a_{jct})$, the share of landlords with assets between the cutoff points a_{jct} and a_{j+1ct}. A similar approach is used by Epple et al. (2020) to define household demand for house quality.

In our case, the household's choice is less straightforward as households choose both tenure and quality subject to a homeownership credit constraint which depends on both income and assets. These additional features make it difficult to identify endowment cut-offs which neatly segment households into different quality bins by tenure. To overcome this issue, I simulate households and aggregate their choices to obtain simulated demand functions by tenure and quality bin. For each unobserved household type i in each city and year, I draw n^{sim} asset and income values (a^l_{ict}, y^l_{ict}) from the joint distribution $F_{ict}(a, y)$. For each simulated household $l = 1, ..., n^{sim}$, I solve the discrete analog of the first order conditions in Equations 4 and 5 to identify each simulated household's optimal quality by tenure:

$$h^r_{ict}(y) = \underset{h_j}{\operatorname{argmax}} \quad u(h_j, y - r_{ct}(h_j) \mid \theta^h_i) \tag{26}$$

$$h^o_{ict}(a, y) = \underset{h_j : p_{ct}(h_j) \leq \overline{p}_{ct}(a,y)}{\operatorname{argmax}} \quad u(h_j, y - c_{ct}(h_j) \mid \theta^h_i) \tag{27}$$

I then compute the tenure shares in Equation 10 based on the optimal quality levels $h^o_{ict}(a^l, y^l)$ and $h^r_{ict}(y^l)$. Finally, I evaluate the integrals in the demand functions in Equations 11–16 by numerically integrating over the simulated households given their housing choices. This gives the simulated demand for each quality bin by tenure.

2.3.2 Parametrizing the Model

We have yet to specify the household's indirect utility $u(h, z \mid \theta^h_i)$ and unobserved user costs of owner-occupation $c_{ct}(h)$. We also need to specify a function to approximate the landlord's variable cost curves over different loan-to-value ratios. I now specify each in turn.

Household Utility from Consumption—Following Epple et al. (2020), I define the indirect utility from living in a house of quality h at cost $r(h)$ given income y as

$$u(h, z \mid \theta_i^h) = \ln\left(1 - \phi_i (h + \eta_i)^{\gamma_i}\right) + \frac{1}{\alpha_i} \ln(y - r(h)) \tag{28}$$

where $\alpha_i, \eta_i, \phi_i > 0$ and $\gamma_i < 0$. The utility function is well-defined when the following constraints are satisfied $1 - \phi_i(h + \eta_i)^{\gamma_i} > 0$ and $y - r(h) > 0$. The estimation imposes both constraints. The parameter vector $\theta_i^h = (\alpha_i, \phi_i, \eta_i, \gamma_i)$ summarizes the taste for housing quality for a household of unobserved type i. The term α_i captures the household's relative preference for housing services compared to non-housing consumption. The term $1 - \phi_i(h + \eta_i)^{\gamma_i}$ may be interpreted as the level of housing services the household generates from housing quality h.

Household Unobserved Usercosts of Owner-Occupation—I specify the household's unobserved user cost of owner-occupation as a linear function in price:

$$c_{ct}(h) = \omega_0 + \omega_1 p_{ct}(h) \tag{29}$$

where ω_0 represents the baseline user cost the household needs to pay for being a homeowner. ω_1 is the amount by which user costs increase for a unit change in prices. It captures the standard components of user costs including depreciation, costs of capital, taxes, and risk premia net of expected capital gains. While the specification is linear in prices, it admits non-linearities in user costs over quality.

The current specification of household user costs makes several simplifications which I intend to improve upon in future work. First, the user cost function is assumed to be the same across all cities. This is restrictive, since local housing markets differ in several ways. On the one hand, London has higher capital gains and lower depreciation rates than other cities, which tend to reduce user costs. On the other hand, because housing is more expensive in London, households may be more likely to rely on borrowing, which increases borrowing costs. Allowing user costs to vary across cities would provide a more realistic description of these differences and help explain the variation in housing demand observed across locations. Second, it would also be useful to allow costs to vary within cities, since households differ in their asset endowments and reliance on borrowing. This would provide a mechanism for capturing differences in demand even within a single city.

While these extensions to the user cost would enrich the demand model, the impact of the current simplification on the main results are likely small. The key goal of the paper is to quantify the equilibrium impact of heterogeneity in landlord supply constraints while controlling for the confounding effects of demand heterogeneity. Any confounding of the landlord supply effects due to the simplified treatment of household user costs is likely to be small given the rich heterogeneity built into the demand model through household preferences and endowments. Household preferences and endowments are allowed to vary by age and city, the same dimensions along which household user costs would be expected to differ. Moreover, the preference structure includes heterogeneity in tastes for both owner-occupation and housing quality. This flexible demand specification should absorb much of the effect of user cost heterogeneity on aggregate demand by quality segment and city.

Landlord's Variable Costs—We need to estimate the landlord's variable costs as a function of the loan-to-value ratio. As discussed in Section 2.2.3, the variable cost plays an important role in the model. When variable costs are convex in the loan-to-value ratio, there is positive sorting between landlords and house quality based on landlord assets. However, it is challenging to directly measure the variable cost since we do not observe many of its components including physical depreciation and the opportunity cost of capital. I instead treat the landlord's cost function as a latent object to be estimated. To approximate the landlord's variable costs net of capital gains, I use the following exponential function with parameters λ_{1c}, λ_{2c} and λ_{3c} which vary by city:

$$vc_{ct}(h \mid a) = \exp\left\{\lambda_{1c} + \lambda_{2c}\left(\lambda_{3c} + 1 - \frac{a}{p_{ct}(h)}\right)\right\} p_{ct}(h) \tag{30}$$

The city-specific parameters allow the function to flexibly capture differences in variable costs across cities. The parameter λ_{1c} influences the baseline level of variable costs which depends on factors including physical depreciation, taxes, risk premia, and expected capital gains. λ_{1c} is expected to be smaller in London consistent with the higher expected capital gains and lower rates of depreciation and risk.[23]

The term $1 - a/p_{ct}(h)$ represents the loan-to-value ratio on a house of quality h for landlords with endowment a. The coefficient λ_{2c} captures the impact of an increase in the loan-to-value ratio on the variable cost. The marginal cost $\partial vc/\partial h$

[23] I impose that the λ parameters do not vary by year or quality segment within each city. In principle, this assumption could be relaxed by measuring the components of λ_{1c} such as capital gains and risk premia outside the model and allowing these estimates to vary by quality segment and year. The estimated components of λ_{1c} could then be introduced as inputs in the model. I intend to pursue this in future work.

is increasing in λ_{2c} when $a/p_{ct}(h) < 1 + \lambda_{3c}$. I impose $a/p_{ct}(h) < 1 + \lambda_{3c}$ in the estimation. When $\lambda_{3c} = 0$, the variable cost is decreasing in a until $a = p_{ct}(h)$ and is constant for higher levels of assets. Positive values for λ_{3c} allows the variable cost to decrease in assets past this point.

The landlord's choice is jointly determined by their participation constraint and the discrete analog of their first order condition. The participation constraint is given by:

$$a \geq \left(1 + \lambda_{3c} + \frac{\lambda_{1c}}{\lambda_{2c}} - \frac{1}{\lambda_{2c}} \ln\left(\frac{r_{ct}(h_j)}{p_{ct}(h_j)}\right)\right) p_{ct}(h_j) \tag{31}$$

where the right hand side of Equation 31 defines the minimum level of assets at which it is profitable for a landlord to operate houses of quality h_j. The participation threshold is increasing in λ_{1c} and λ_{3c}. The participation threshold is increasing in λ_{2c} whenever the rent is greater than $\exp\{\lambda_{1c}\}p_{ct}(h_j)$ which represents the lowest possible cost a landlord can face. This condition is always satisfied as landlords do not operate below the break even point.

As discussed in Section 2.3.2, the asset value which satisfies the indifference condition $\Pi_{ct}(h_j \mid a) = \Pi_{ct}(h_{j+1} \mid a)$ represents the cut-off a_{jct} at which the landlord is indifferent between renting out houses of quality h_j and h_{j+1}. Given the definition of variable costs in Equation 30, the indifference condition becomes:

$$r_{ct}(h_{j+1}) - r_{ct}(h_j) = \exp\left\{\lambda_{1c} + \lambda_{2c}\left(\lambda_{3c} + 1 - \frac{a}{p_{ct}(h_{j+1})}\right)\right\} p_{ct}(h_{j+1}) \; - $$
$$\exp\left\{\lambda_{1c} + \lambda_{2c}\left(\lambda_{3c} + 1 - \frac{a}{p_{ct}(h_j)}\right)\right\} p_{ct}(h_j) \qquad (32)$$

The solution to this indifference condition represents the asset cut-off a_{jct}. It is not possible to explicitly solve for the asset value which satisfies Equation 32. However, we can solve for it numerically provided a unique solution exists. To verify that a unique solution does exist, consider the right hand side of Equation 32. We know that it intersects with the positive constant $r_{ct}(h_{j+1}) - r_{ct}(h_j)$ at most once because it is strictly decreasing in a.[24] We also know that it intersects with $r_{ct}(h_{j+1}) - r_{ct}(h_j)$ at least once because it spans the entire non-negative range from zero to ∞. This verifies that a solution to Equation 32 exists and is unique.

Evolution of the Housing Stock—Let m and m' index two different markets. Similar to Epple et al. (2020), I specify the evolution of the housing stock by quality segment by the following constant elasticity function:

$$\frac{N_{m'}(h_j)}{N_m(h_j)} = \nu_{m'}\left(\frac{\kappa_{m'}}{\kappa_m}\right)^{h_j \times \zeta_{pop}} \qquad (33)$$

[24]This difference on the right hand side is the discrete analog of the derivative of variable costs with respect to quality. We know this difference is increasing in $\lambda_{3c} + 1 - a/p_{ct}(h_j)$ due to the convexity of the variable cost function. Since $\lambda_{3c} + 1 - a/p_{ct}(h_j)$ is decreasing in a, it follows that the difference in the right hand side is also strictly decreasing in a.

where $\nu_{m'}$ is a constant to ensure that the housing stock in m' sums to one. Equation 33 implies that a one percent increase in the relative population density $\kappa_{m'}/\kappa_m$ generates a $h_j \times \zeta_{pop}$ percent change in $N_{m'}(h_j)/N_m(h_j)$, the housing stock in m' relative to m. When $\zeta_{pop} > 0$, the relative supply of high quality houses is greater in denser cities. In contrast, when $\zeta_{pop} < 0$, the relative supply of high quality houses is greater in less dense cities.

2.3.3 Approximating the Empirical Distribution of Endowments, Prices and Rents

Endowments, prices, and rents are important inputs in the structural estimation. I model the marginal distributions of prices, rents, household income and household and landlord assets using a generalized beta distribution of the second kind, henceforth referred to as GB2. The GB2, also known as the generalized beta prime distribution, is a flexible four parameter distribution which nests many of the distributions commonly used to model income, assets and prices including the log normal, Singh-Maddala, and generalized gamma distributions (McDonald, 1984; McDonald and Xu, 1995). I use a Student-t copula to model the dependence between household income and assets. The joint distribution of household endowments is allowed to vary by market and demographic group, while the distribution of landlord assets, as well as prices and rents, varies by market. The estimation is conducted using Maximum Likelihood.

2.3.4 Solving for the Equilibrium

Since house quality is latent, I start by normalizing the J levels of house quality based on the house price distribution in a base market. I define the base market to be London in 2014. I divide the house price distribution in the base market into J bins. I then define quality h_j to equal the price at the j^{th} cut-off. Given the normalization the equilibrium price simply equals the quality, $p_{00}(h_j) = h_j$ in the base market. While prices are linear in quality in the base market due to the normalization, it need not be so in other markets. We can also directly measure the equilibrium rent $r_{00}(h_j)$, overall housing stock $N_{00}(h_j)$, stock of rentals $N^r_{00}(h_j)$, and stock of owner-occupied houses $N^o_{00}(h_j)$ in each quality segment of the base market. We therefore do not need to solve for the equilibrium in the base market.

We need to solve for the equilibrium in all other markets. Given the normalization of the base market, we can use the recursion in Equation 33 to compute the housing stock in London in 2015 and in all other cities in 2014. Suppose we have computed the housing stock based on this recursion for some market (c, t). We can then solve for the vector of rents and prices which satisfy the equilibrium conditions in Equations 24 and 25 for market (c, t). Repeating this process for subsequent years yields the housing stock and the equilibrium in all markets.

2.3.5 Method of Moments Estimator

I estimate the model using a method of moments estimator.

Parameters—We need to estimate three sets of parameters which shape the choices of households, landlords, and the builder. First, households belong to type i with probability $p_{i|x}$ and their preferences over housing quality and tenure depend on the parameters $\theta_i^h = (\eta_i, \phi_i, \gamma_i, \alpha_i)$ and $\theta_i^\tau = (\beta_i^{own}, \sigma_{logit})$ respectively. Moreover households' user costs of owner-occupation depend on $\omega = (\omega_0, \omega_1)$ and they face credit constraints which depend on the mortgage thresholds $\phi = (\phi^{lti}, \phi^{ltv})$. In the estimation, I fix $\phi^{ltv} = 0.9$ and only estimate ϕ^{lti}. Second, the landlord's profit function depends on variable cost parameters $\lambda = (\lambda_{1c}, \lambda_{2c}, \lambda_{3c})$. Lastly, the evolution of the housing stock depends on the elasticity ζ_{pop}. Let $\Omega = (p_{i|x}, \theta_i^h, \theta_i^\tau, \omega, \phi, \lambda, \zeta_{pop})$ be the collection of all the parameters that need to be estimated.

Moments—The estimation uses the following set of moments for each city and year: (i) marginal distributions of prices and rents; (ii) distribution of rent-to-price ratio over price;[25] (iii) correlation between income and prices and rents by age group; (iv) assets-to-house price ratio for owner-occupiers by age group; (v) ownership rates by age group; and (vi) mapping of ranks in the rent distribution to ranks in the price distribution.[26] These moments characterize the sorting of households and landlords into tenure and quality segments in equilibrium across cities and over time.

[25]In each iteration of the estimation, the model outputs a rent $r_{cy}(h)$ and a price $p_{cy}(h)$ by house quality h. The rent-to-price moment condition for quality h in cy is defined as the difference between the model-based ratio $r_{cy}(h)/p_{cy}(h)$ and the data counterpart for the rent-to-price ratio. I impute the data-based rent-to-price ratio using mappings of the type discussed in Section 1.2.1. In particular, using the mappings I impute the price in cy given rent $r_{cy}(h)$ from the model. Dividing $r_{cy}(h)$ by the imputed price gives the rent-to-price ratio implied by the data i.e., the data-based moment.

[26]Moments (ii) and (vi) are based on mappings of the type discussed in Section 1.2.1. In Section 1.2.1, these mappings are estimated using polynomial regressions with the polynomial order selected by cross-validation. To construct the moments for the structural estimation, I instead use mappings of the same type estimated using third-degree polynomials. The impact of this simplification should be small, since the mappings based on the third order polynomial are similar to those obtained in Section 1.2.1.

Estimation Algorithm—The estimation imposes market clearing for each quality segment and each type of tenure in each market. The method of moments estimator selects parameters which minimize the distance between equilibrium prices, rents and allocations in the model and in the data across all cities and years. The algorithm is implemented as follows:

1. Initialize the parameter vector to Ω_0
2. Draw $N^{sim} = 5000$ households per unobserved type per market. The households are drawn from $F_{ict}(a, y)$, the joint distribution of assets and income for unobserved type i. Recall from Equation 2 that $F_{ict}(a, y)$ depends on the type probability $p_{i|x}$.
3. Normalize quality using prices in London in 2014. Directly measure the housing stock and the schedule of prices and rents by quality in the base market.
4. Calculate excess demand in the base market given the normalization:

    (a) Given the normalized quality, prices and rents, compute the simulated demand by tenure and quality segment, $D^r_{ct}(h, p_{ct}, r_{ct})$, $D^o_{ct}(h, p_{ct}, r_{ct})$, and the supply of landlord services by quality segment, $S_{ct}(h, p_{ct}, r_{ct})$.

    (b) Check if the equilibrium conditions in Equations 24 and 25 are satisfied. The base market will not clear for an arbitrary guess of Ω. Define $\mathrm{ED}_{00}(\Omega)$ as the excess demand in the base market which needs to be minimized.

5. Solve for the equilibrium in all markets other than the base market (London in 2014):

(a) Initialize the vector of prices p_{ct} and rents r_{ct}.

(b) Compute the housing stock by quality segment recursively

(c) Compute the simulated household demand by tenure and quality segment

(d) Compute the supply of rental services by landlords by quality segment

(e) Update the price and rent vectors and repeat steps 5b–5d until the equilibrium conditions are satisfied for the market in question

6. Form moment conditions based on differences between the model-based moments and their empirical counterparts. Construct an objective function $M(\Omega)$ based on the moment conditions and excess demand in the base market $\mathrm{ED}_{00}(\Omega)$.

7. Update Ω and repeat steps 2–6 until $M(\Omega)$ is minimized.

2.3.6 Identification

Separately Identifying the Parameters for Households, Landlords, and Builders—We can directly measure the housing stock in the base market given our normalization which equates house quality to prices in the base market. Since the baseline stock is normalized outside the structural model, the structural estimates of the housing supply parameter ζ do not impact the equilibrium in the base market. This exclusion restriction helps identify the parameters for the household and landlord separately from that of the builder.

I use two main sources of identification to distinguish household preferences from landlord cost parameters. First variation in household and landlord endowment distributions across markets helps separately identify the parameters for households and landlords. The model provides a mapping of household and landlord endowment distributions to equilibrium outcomes, with household parameters moderating the effect of household endowments on the equilibrium and landlord parameters shaping the effect of landlord endowments. In the estimation I exploit independent cross-market variation in the endowment distributions of households and landlords to identify household and landlord parameters.[27] Second, the estimation uses several demand-side moments that do not depend on landlord parameters conditional on prices and rents. These include ownership rates by age, correlations between house quality and household endowment by age, and quality segmentation by tenure. These exclusion restrictions further help me disentangle household preferences from landlord costs.

In the estimation, I allow household unobserved type probabilities and landlord cost parameters to vary across cities but hold them constant over time. This means that differences in equilibrium outcomes between cities can arise either due to cross-city differences in these unobserved parameters or due to heterogeneity in the distribution of observed endowments. To disentangle these two effects, I use data from three different years for each city. Given the restriction that household and landlord parameters do not vary over time, the data from multiple years allows me to consider the effect of varying endowments while holding unobserved costs and

[27]This approach closely relates to the literature on identifying hedonic models using data from multiple markets (see e.g., Heckman et al., 2010).

preferences fixed.

Identifying the Household's Parameters—The parameter governing the household's taste for tenure β_i^{own} and the idiosyncratic shock $\varepsilon_{i\tau}$ drop out of the first order conditions in Equations 4 and 5 as they do not affect the choice of quality conditional on tenure. This exclusion restriction helps identify the taste for tenure parameters separately from the user costs of owner-occupation which impact both tenure choice as well as differences in the choice of quality conditional on tenure. Additive separability of the household's taste for tenure is not necessary to generate this exclusion restriction. The exclusion restriction arises even when the taste for tenure is introduced multiplicatively in Equation 3.

The relative distribution of household endowments and house prices varies by household characteristic x, city, and time. The credit constraint parameters, ϕ^{ltv} and ϕ^{lti}, influence how ownership rates respond to variation in the distribution of household endowments relative to local house prices. This correlation helps identify the credit constraints.

The type probabilities $p_{i|x}$ are identified by two forces. First, differences in the type probabilities by observed characteristic x generate differences in aggregate demand by x. Second, as households with characteristic x become more homogenous (i.e., as $p_{i|x}$ converges to one for any given i), the correlation between income and rents converges to one for households with x. In the data the correlations are substantially lower than one implying the existence of household heterogeneity. These two forces discipline the estimates for $p_{i|x}$.

2.4 Empirical Results

2.4.1 Parameter Estimates

This section presents the parameter estimates for the household, landlord, and housing supply.

Household Preference Parameters—

Table 1: Household Preference Parameters

	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6
α	18.538	5.11	4.596	1.94	0.908	3.687
η	7.447	3.013	1.75	0.652	0.618	2.456
γ	-1.349	-3.037	-2.548	-3.644	-1.377	-3.621
ϕ	5.385	5.047	0.593	6.182	2.899	19.779
$\ln \beta_{logit}$	1.488	0.495	-10.858	1.425	2.867	11.795
σ_{logit}	0.303	0.303	0.303	0.303	0.303	0.303

Note: This table presents estimates for the household preference parameters

The model is estimated with six unobserved household types, with type probabilities varying by age and by whether the household resides in London. Table 1 presents the utility parameters for each unobserved type, while Figure 10 plots the distribution over types by household age, both within and outside London. The types are ordered by the share of young households belonging to the type. Figure 10 reveals a strong negative correlation in type probabilities across age groups. Whereas young households are most likely to be type 1 and least likely to be type 6, the opposite is true for older households. Moreover, the type probabilities exhibit substantial heterogeneity within each age group, with over 70% of households in either age group belonging to one of two types. These patterns indicate that the demand system effectively captures rich heterogeneity in preferences among households conditional on income and assets.

The estimates for α generally decrease with the type, suggesting that younger households have a higher relative preference for housing consumption compared to non-housing consumption. Similarly, the estimates for η tend to decrease with the type, indicating that younger households derive greater utility from lower levels of housing compared to their older counterparts. The estimates for γ and ϕ also exhibit heterogeneity, however, they do not vary as systematically with the age-specific type probabilities.

Figure 10: Conditional Distribution of Households Over Unobserved Types

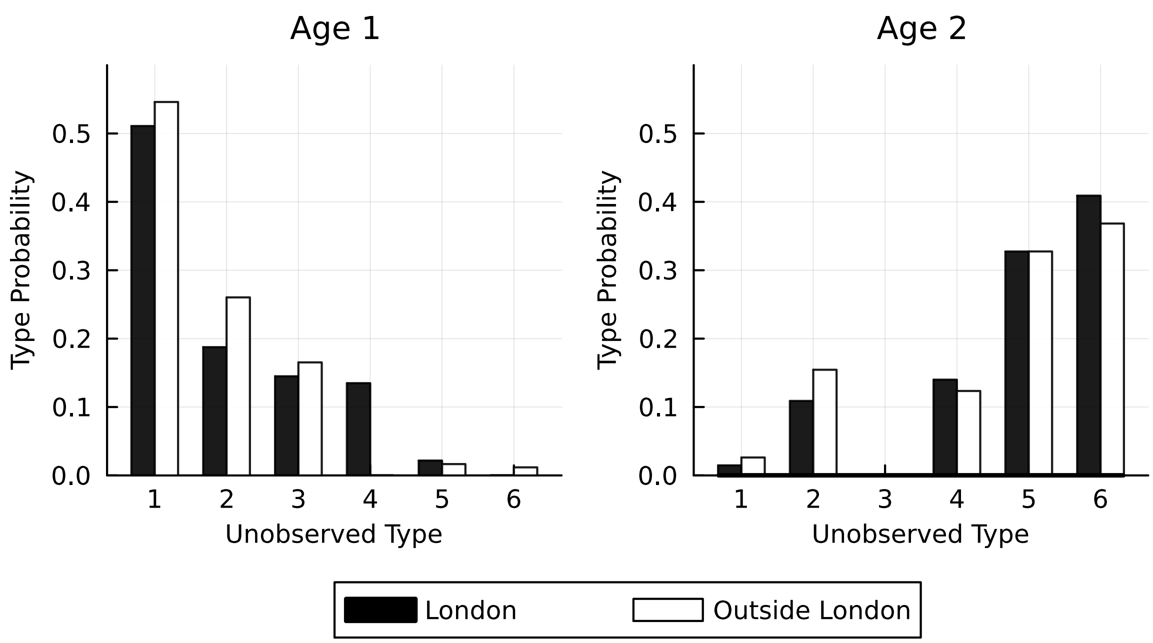

Note: This figure plots estimates for the share of households who belong to each unobserved type by age and location

Table 2: Landlord Cost Parameters

	London	Manchester	Birmingham
λ_1	-5.361	-4.667	-4.435
λ_2	2.112	0.821	0.348
λ_3	0.901	3.481	4.321

Note: This table presents estimates of the landlord's cost parameters by city

Turning to households' deterministic preferences for tenure, $\ln \beta_{logit}$, Table 1 indicates that all types except type 3 prefer owner-occupation to renting. Type 3 households have a large and negative estimate for $\ln \beta_{logit}$ implying a strong preference for renting. Approximately 15% of younger households fall into this category, while none of the older households do, suggesting that all older households prefer owning over renting. Among households that prefer owning, the ownership premia tends to be significantly larger for older households. The estimate for the scale parameter σ_{logit} is smaller in magnitude than the deterministic ownership preference across all types.

Mortgage Thresholds— The loan-to-income threshold is estimated to be 5.469, implying that households can borrow up to approximately 5.5 times their annual income. The estimation fixes the loan-to-value ratio to 0.9 which implies that households need to make a minimum deposit of 10% on the value of the house.

Unobserved Usercosts of Owner-Occupation—The estimated parameters for the user cost functions, ω_0 and ω_1, are 0.446 and 0.029 respectively. The intercept ω_0 is ignorable considering that prices are in the hundred thousands. This implies that the user-cost-to-price ratio is effectively constant and equals 0.029.

Landlord's Parameters— Table 2 presents estimates for the parameters of the landlord's cost function by city. As expected, the estimate for λ_{1c} shows that baseline costs are lowest for landlords in London and highest for landlords in Birmingham. This result is consistent with the higher realized capital gains and the lower depreciation rates and risk in London.[28] Conversely, the estimate for λ_{2c} suggests that landlord costs are most sensitive to borrowing in London and least sensitive in Birmingham. The estimate for λ_{3c} is greatest in Birmingham. Higher values of λ_{3c} tend to make the participation constraint more binding.

Housing Supply Parameters—The estimate for the housing supply parameter ζ_{pop} is -1.259 indicating that the quality distribution of the housing stock is strongly and negatively related to the density or congestion in that market. Higher quality housing is relatively less abundant in denser cities such as London.

2.4.2 Model Fit

This section presents results which illustrate the model's fit. In the interest of space, I only present comparisons for 2015. The fit is similar in other years.

Market Clearing—The model successfully achieves market clearing. Appendix Figure B.4 plots the cumulative demand and supply functions for the overall housing market and for the rental market in 2015. The figure illustrates that demand very closely corresponds to supply over the range of quality in all cities.

[28]See Chapter 1 for evidence and discussion of these differences across cities

Selected Targeted Moments—Figure 11 compares the pattern of quality segmentation by tenure generated by the model to the pattern of segmentation estimated directly from the data. The fit is very good. The model is able to replicate differences in the patterns of segmentation across cities with a high degree of accuracy.

Figure 11: Quality Segmentation in Rental vs. Owner-Occupied Markets Across Cities

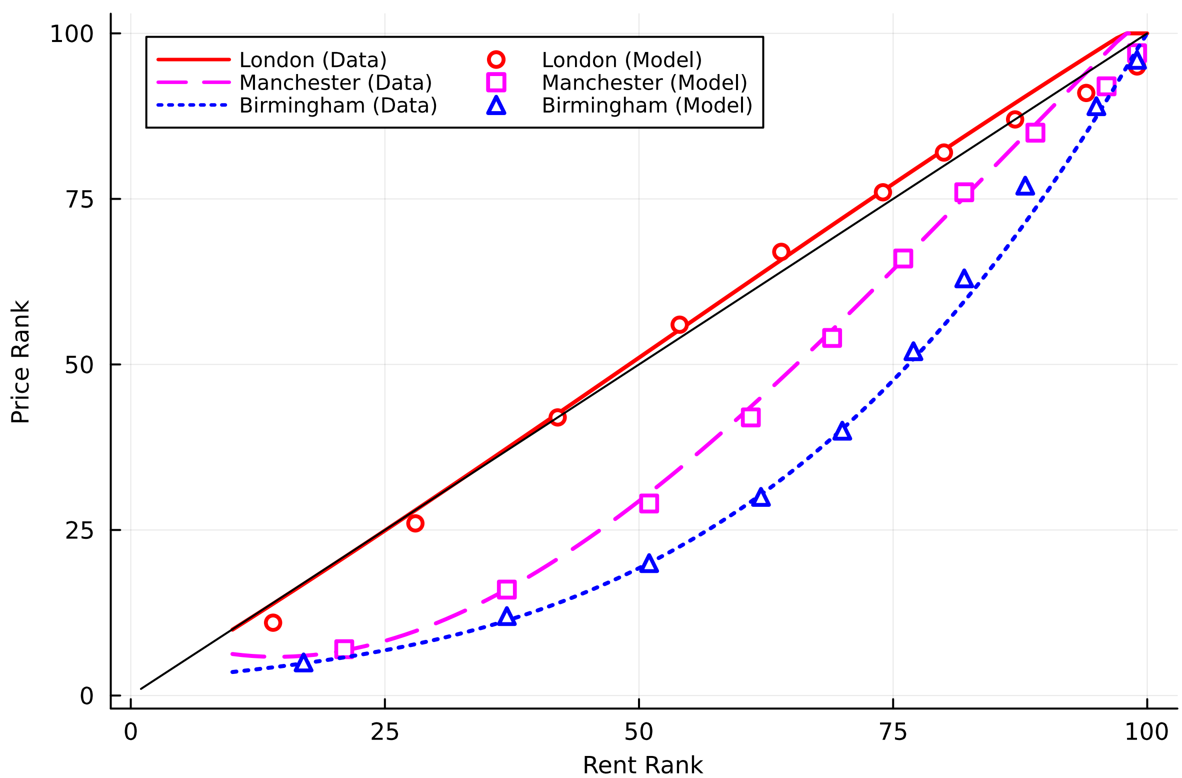

Note: This figure compares the relative distribution of quality between the rental and owner-occupied segments of the housing market across cities. The lines represent the relationship observed in the data for each city, while the scatter points show the relationship as predicted by the estimated equilibrium model.

Appendix Figure B.5 plots the cumulative distribution of prices and rents for each city. The model successfully generates the empirical price and rent distributions. Appendix Figure B.3 compares ownership rates predicted by the model to rates obtained from the data. The model performs fairly well for both age groups in Manchester and Birmingham. In London, it correctly generates a higher ownership rate for older households as observed in the data, however, it understates the magnitude of this difference.

Landlord Assignment to Quality—The model predicts that landlords sort positively to quality based on their asset endowments. In the discretized version of the model with ten quality levels, the assignment is determined by nine asset cutoffs which satisfy the indifference condition in Equation 32 for each quality level. Appendix Figure B.6 plots the asset cutoffs for each quality level by city. As expected, the asset cutoffs are monotonically increasing in house quality.

Figure 12 presents the profit curves for landlords at each of the nine asset cutoffs. Consistent with the model's prediction, the profit curves are flat at the indifference point corresponding to each asset cutoff. This pattern illustrates how the assignment model partitions the landlord asset space into different quality levels based on profit maximization. The minimum profit for landlords operating within each quality level is represented by the red step function defined by the flat sections of each of the nine profit curves in the figure. This step function indicates that realized profits increase with the quality of the property and landlord assets.

The results presented in this section indicate that the estimated model adequately captures the key forces which determine the joint equilibrium in the rental and owner-occupied sectors of the housing market. In the following sections, I use the estimated model to analyze the role of landlord heterogeneity in generating patterns observed both within and across housing markets.

2.5 Counterfactuals

The model captures landlord heterogeneity in two ways—through variations in observed wealth within cities and through differences in unobserved opportunity costs of capital both within and across cities. In this section I conduct three counterfactual simulations to quantify the impact of this heterogeneity on housing markets across cities. I find three key insights from these simulations.

First, in Section 2.5.1, I show that within-city variation in landlord wealth almost completely explains the downward slope of the rent-to-price function with respect to property prices within cities. Second, in Section 2.5.2, I show that cross-city differences in landlords' opportunity costs significantly drive the observed patterns of quality segmentation and rent-to-price ratios across cities. Moreover I show that these differences have important implications for the distribution of household welfare. Finally, in Section 2.5.3, I illustrate how the impact of landlord taxes on rents and prices varies across cities and quality segments due to differences in landlord composition.

Figure 12: Estimated Landlord Profit Curves by Landlord Asset Endowment (2015)

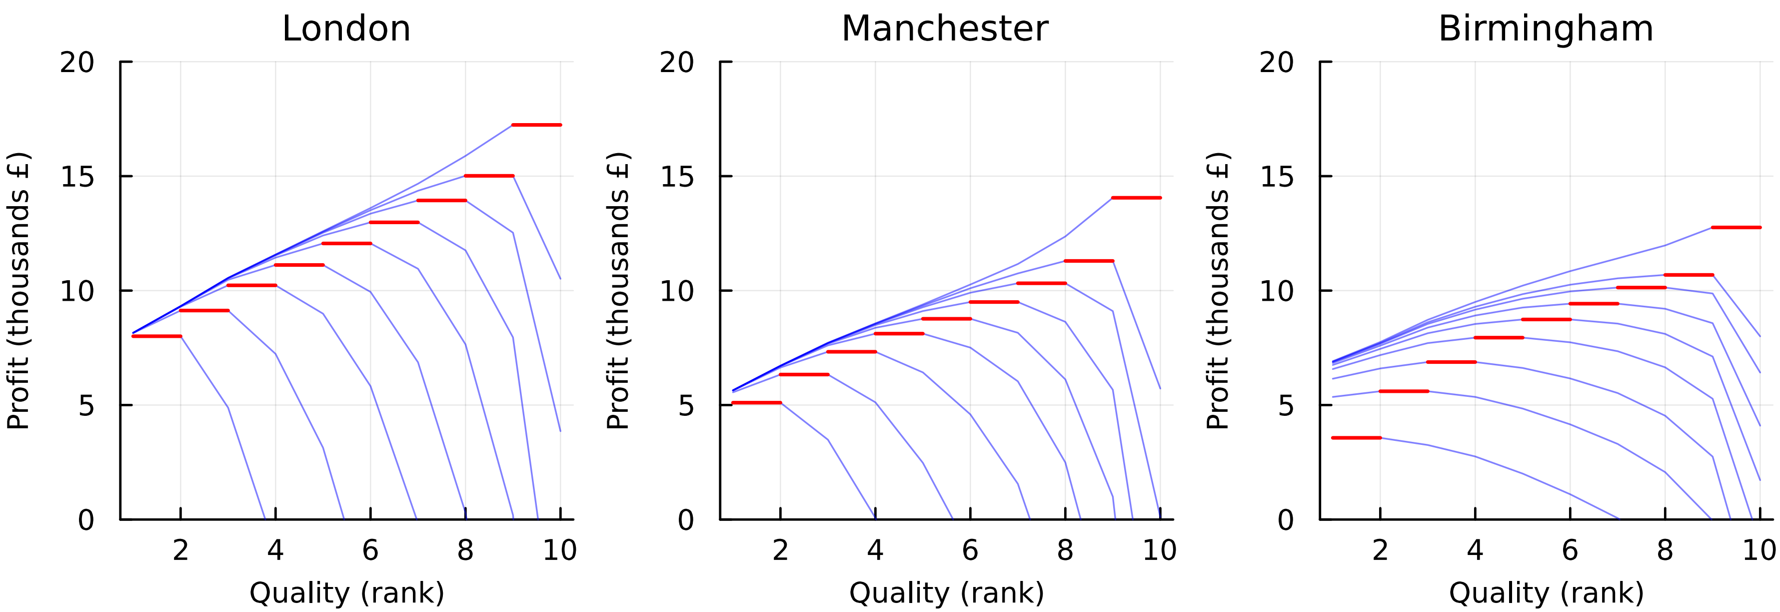

Note: This figure shows estimated profit curves for landlords at nine different asset cutoffs. The topmost curve represents profits for landlords with assets which make them indifferent between the highest and second-highest quality levels. The bottommost curve represents profits for landlords with assets which make them indifferent between the lowest and second-lowest quality levels. The indifference point of each curve is highlighted red.

The counterfactuals hold the distribution of housing quality fixed. This assumption is appropriate for the short run, before builders have time to adjust the composition of the housing stock. However, in the long run, changes in relative rents and prices would likely alter builder incentives to supply properties of different quality. By abstracting from this response, the counterfactuals may tend to overstate price and rent adjustments to shocks in landlord costs or wealth. This limitation is likely less important in cities like London, where builder supply constraints are greater. I therefore interpret the counterfactuals as capturing a short-run response. In future work, I plan to extend the model to allow the housing stock to adjust endogenously to changes in rents and prices.

2.5.1 Impact of Within-City Heterogeneity in Landlord Endowments

Several studies have found that the rent-to-price function decreases in price (Verbrugge, 2008; Verbrugge and Poole, 2010; Heston and Nakamura, 2011; Bracke, 2015; Halket et al., 2020) both in the UK and US. The literature has proposed several explanations including differences in rental vacancy rates and unobserved building characteristics. In this section I show that heterogeneity in landlord wealth endowments naturally generate a downward sloping rent-to-price function in equilibrium, even when one controls for unobserved quality. Moreover, using counterfactual simulations I show that the downward slope of the rent-to-price function can be eliminated by reducing the heterogeneity in landlord wealth endowments.

Figure 13: Demand and Supply of Rental Properties When the Rent-to-Price Ratio is Fixed (London 2015)

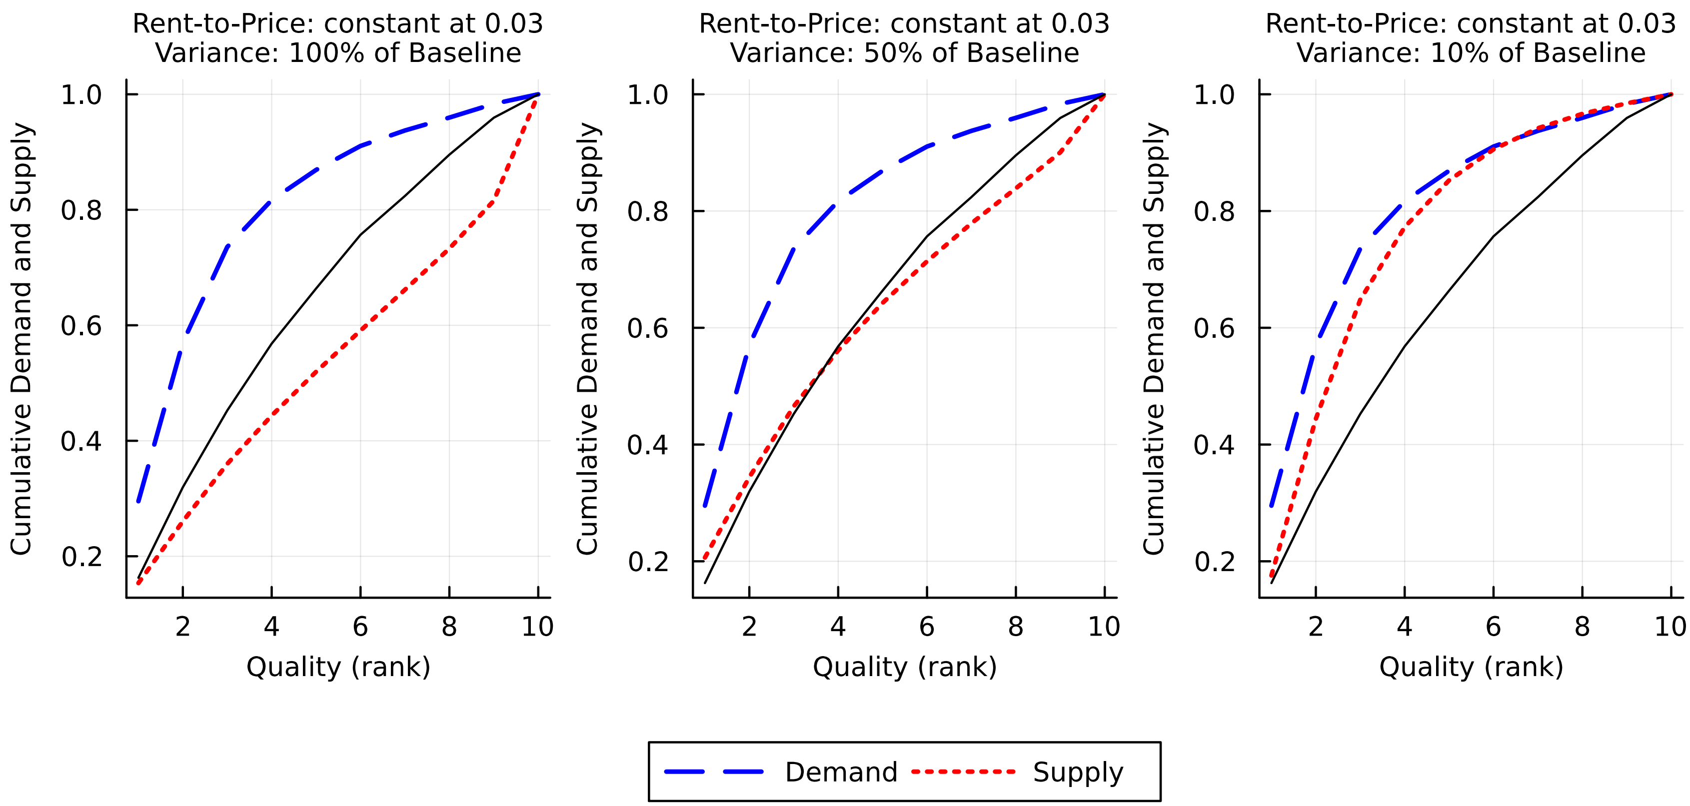

Note: This figure presents cumulative demand (blue dashed curve) and cumulative supply (red dotted curve) when the rent-to-price ratio is fixed to equal 0.03. For reference, the figure also plots the baseline equilibrium cumulative quantities when the rent-to-price ratio is allowed to vary (solid black curve). The center and right-most plots exogenously reduce the variance of the landlord's asset distribution by 50% and 90% respectively. The left-most plot leaves the landlord's asset distribution unchanged. The gap between the demand and supply curves represents the extent by which the rental market fails to clear. To fix the rent-to-price at 0.03, I set annual rents to equal three percent of the equilibrium price at baseline.

Figure 13 plots the cumulative demand and supply of rental services over quality in London when the rent-to-price function is exogenously fixed to equal 0.03 throughout the quality distribution. To fix the rent-to-price at 0.03, I set annual rents to equal three percent of the equilibrium price at baseline. When there is excess supply in a quality segment, the cumulative supply function is steeper than the cumulative demand function. The opposite is true when there is excess demand. The leftmost plot shows that fixing the rent-to-price to 0.03 generates excess demand at lower levels of quality and excess supply at higher levels. The market does not clear because when the rent-to-price is constant, landlords are more attracted to higher levels of quality than households. For the market to clear the attractiveness of higher quality must fall for landlords. This is achieved when the slope of the rent-to-price function with respect to price becomes more negative. This illustrates how the negative slope of the rent-to-price function arises due to market clearing forces.

To explore why the market does not clear when the rent-to-price is constant, the center and rightmost subfigures of Figure 13 plot demand and supply for rental services with the variance of the landlord asset distribution reduced by 50% and 90% respectively. The market clearing condition comes close to being satisfied when the variance of the landlord endowment distribution is reduced by 90%. These plots suggest that the rent-to-price function is steeper when landlord assets are more spread out. This can occur in the model because landlords sort positively to quality based on wealth which implies that the increased variance increases landlord wealth in the higher quality segments and reduces it in the lower segments.

To directly test whether reducing the variance of landlord endowments indeed flattens the rent-to-price function, Figure 14 plots counterfactual rent-to-price functions obtained by resolving the housing market equilibrium under different landlord asset distributions. The figure on the left plots the rent-to-price function in London with the mean landlord wealth set to different levels but with the variance fixed to the baseline level. The figure on the right plots analogous rent-to-price functions with the variance reduced to 10% of the baseline variance. The differences are stark. The reduction in variance almost completely flattens the rent-to-price function regardless of the mean level of wealth.[29] These results illustrate the significant influence of landlord heterogeneity on the shape of the rent-to-price function over quality segments within cities.

2.5.2 Impact of Cross-City Heterogeneity in Landlord Costs

The estimates of the landlord's cost function in Table 2 show that at low levels of borrowing, the opportunity cost for landlords is significantly lower in London than in Birmingham.[30] I show in this section that these cross-city differences in unobserved landlord costs help explain much of the observed differences in housing markets across cities. Moreover, these cost differences have important consequences for the distribution of household welfare.

[29] The mean of the wealth distribution tends to influence the level of the rent-to-price function.

[30] For instance when landlord wealth exactly equals the house price, the opportunity cost is 40% lower in London than in Birmingham. The difference in costs is increasing in the asset to price ratio.

Figure 14: Equilibrium Rent-to-Price Functions Under
Counterfactual Landlord Asset Distributions (London 2015)

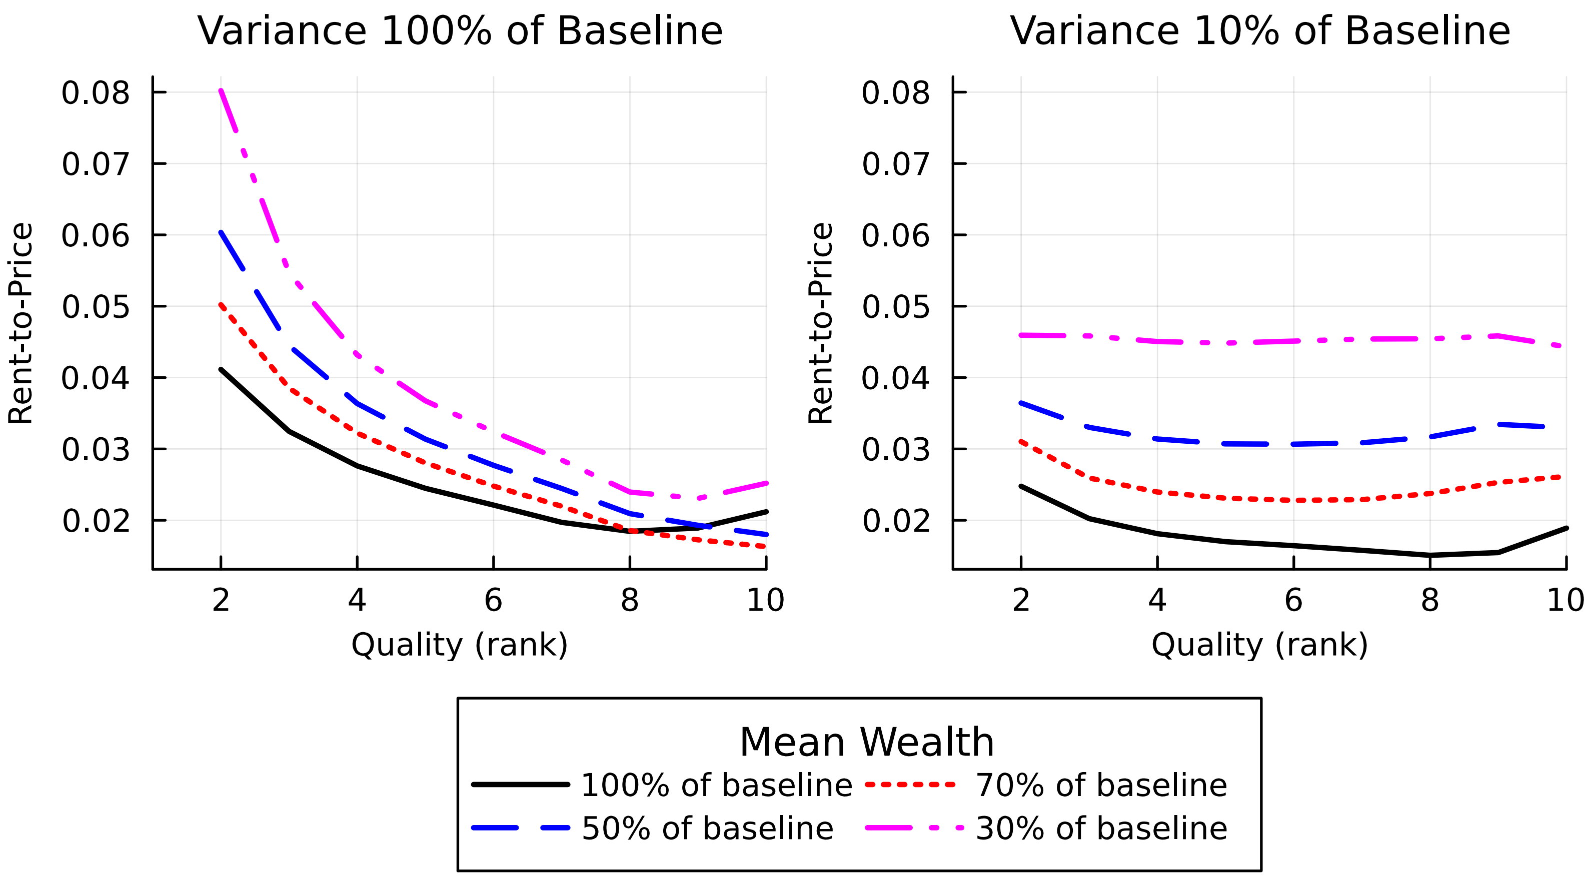

Note: This figure presents equilibrium rent-to-price functions for different landlord asset endowments. The figure on the left plots the rent-to-price function holding the variance fixed and varying the mean of the landlord wealth distribution. The figure on the right presents analogous rent-to-price functions with the variance reduced by 90%.

To quantify the impact of cross-city differences in landlord costs, I simulate a counterfactual in which I fix the cost function in London to equal the cost function in Birmingham and solve for the resulting equilibrium in London. Figure 15 compares equilibrium outcomes in the counterfactual scenario in London against the baseline equilibrium in London and Birmingham. Figures 15a and 15b show that the change in landlord costs causes rents to increase relative to both prices and user costs of owner-occupation in London.[31] This is consistent with landlord costs being higher in Birmingham in baseline.

Comparing the gap between London and Birmingham in the counterfactual and baseline scenarios shows that the cross-city difference in landlord costs accounts for much of the differences in housing markets across cities. In particular, the cross-city difference in landlord costs explains up to 40% of the gap in the rent-to-price ratio between the two cities and up to 67% of the gap in the ratio of user costs of owner-occupation and rents. These results show that the cost advantage faced by landlords in London generates important differences in rental supply between the two cities which in turn impacts the overall housing market in each city. Due to the cost advantage, landlords in London are willing to accept significantly lower rents while paying much higher prices to acquire properties. These differences in rental supply have large effects on the allocation of housing quality across the rental and owner-occupied sectors. Figure 15c shows that equalizing landlord costs

[31]The sensitivity of the rent-to-price and user cost-to-rent ratios to landlord costs is increasing in the quality segment. This occurs because in the model an increase in costs reduces the optimal quality for the marginal landlord in each segment. The supply effect of a cost increase is less pronounced in lower quality segments because marginal landlords in these segments can more easily be replaced by landlords who were previously operating in higher segments. This substitution effect is less pronounced in higher quality segments.

Figure 15: Impact on Rents, Prices, and Allocations
Counterfactual: Increasing Landlord Costs in London to Birmingham Levels

(a) Rent-to-Price Ratio

(b) User cost-to-rent ratio

(c) Ratio of Segmentation

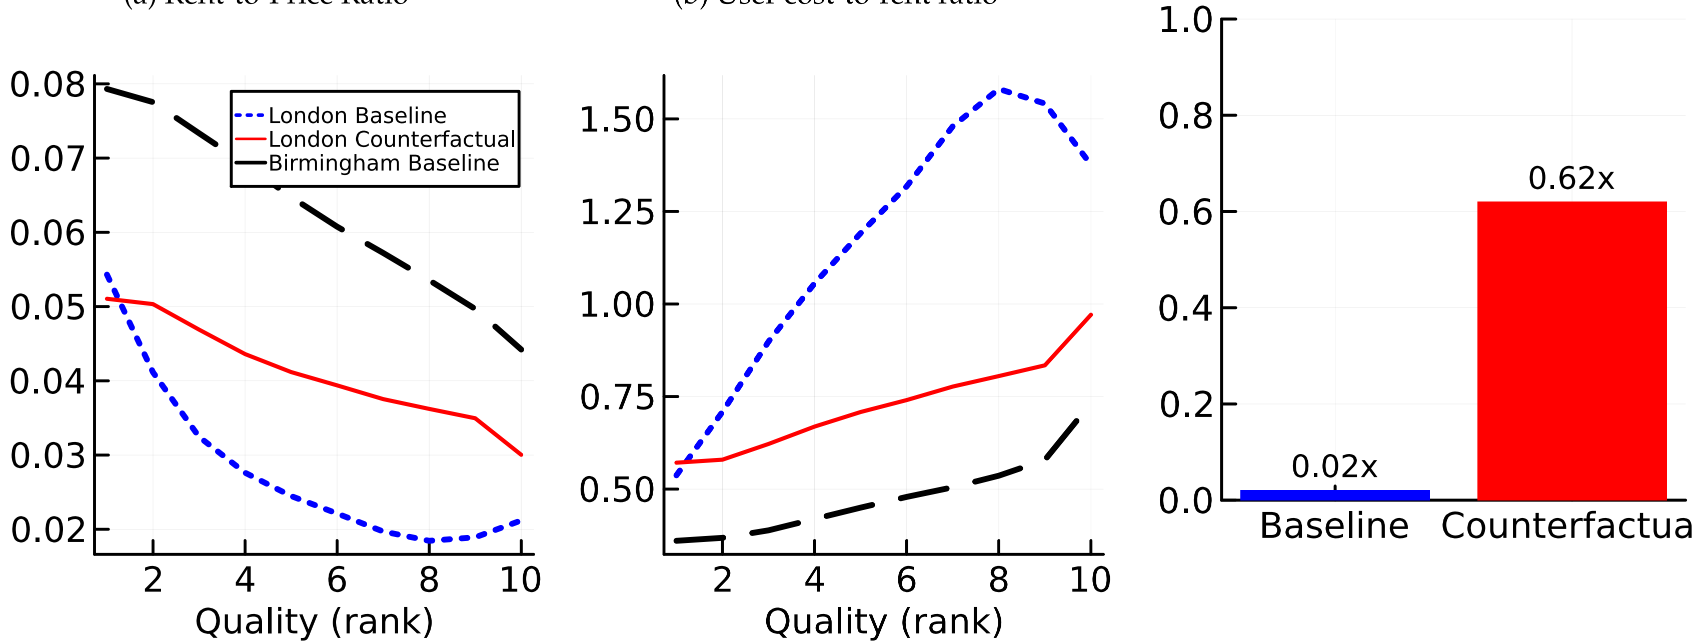

Note: This figure compares equilibrium outcomes between the counterfactual and baseline scenarios. Figure 15a compares the counterfactual rent-to-price function in London (in red) against the baseline rent-to-price functions in London (in blue) and Birmingham (in black). Similarly, Figure 15b compares the counterfactual ratio of user costs-to-rents in London (in red) against the baseline in London (in blue) and Birmingham (in black). Figure 15c compares the ratio of segmentation between London and Birmingham in the counterfactual (in red) and baseline (in blue) scenarios.

between London and Birmingham reduces the cross-city gap in quality segmentation between the rental and owner-occupied sectors by 60%.

Welfare Effects of Heterogeneity in Landlords Costs—The results so far show that cross-city heterogeneity in unobserved landlord costs generate large cross-city differences in the allocation and costs of housing across the rental and owner-occupied sectors. I now explore the welfare consequences of these differences. Figure 16 plots the distribution of compensating variation for households in London when landlord costs in London are increased to Birmingham levels. The distributions are presented separately by the tenure of the household at baseline. The plots show that the welfare effects vary greatly by tenure. The cost advantage faced by landlords in London increase the welfare of the median renter household in London by £1,448 per year (3.6% of the median household income in London). This reflects the presence of lower costs and higher quality in the rental sector of London as shown in Figure 15. On the other hand, the median owner-occupier is worse-off by £2,139 per year (5.3% of the median household income in London), reflecting the higher costs and worse relative quality in the owner-occupied sector of London.

Figure 16 shows that there is large heterogeneity in the welfare effects even conditional on tenure. To explore the source of this heterogeneity, Figure 17 plots the distribution of compensating variations by household income quartile. The figure shows that the magnitude of the welfare effects are increasing in income for both renters and owner-occupiers. This occurs because in the model a reduction in landlord costs increases the profit-maximizing quality for marginal landlords in each

Figure 16: Distribution of Welfare Effects by Tenure
Counterfactual: Increasing Landlord Costs in London to Birmingham Levels

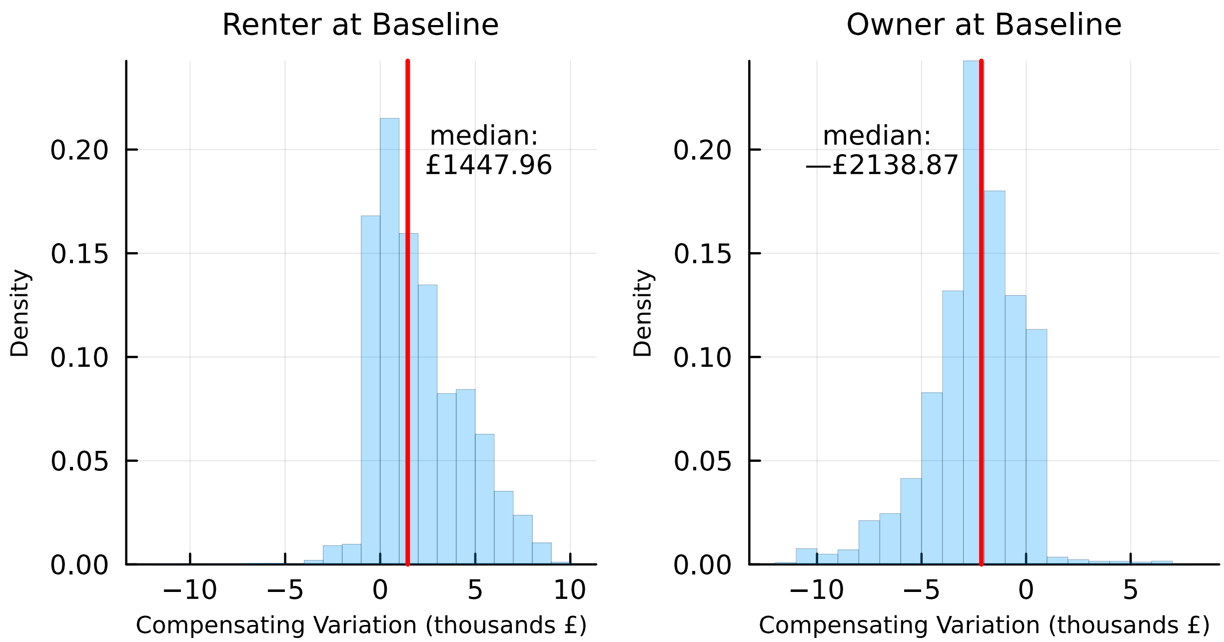

Note: This figure plots the distribution of compensating variation for households associated with increasing landlord costs in London to Birmingham levels. The plot on the left presents the compensating variation for households who rented at baseline, whereas corresponds to households who owner-occupied at baseline.

quality segment. This tends to increase the relative supply of rental services in higher quality segments in turn magnifying the effect on rents and prices in those segments.

2.5.3 Influence of Landlord Heterogeneity on the Impact of Housing Policy

The previous sections show that landlord heterogeneity in wealth and costs play a crucial role in shaping housing markets both within and across cities. In this section I examine how the equilibrium effects of housing policy are influenced by the composition of landlords across local housing markets. Specifically, I explore the

Figure 17: Heterogeneous Welfare Effects by Household Income and Tenure
Counterfactual: Increasing Landlord Costs in London to Birmingham Levels

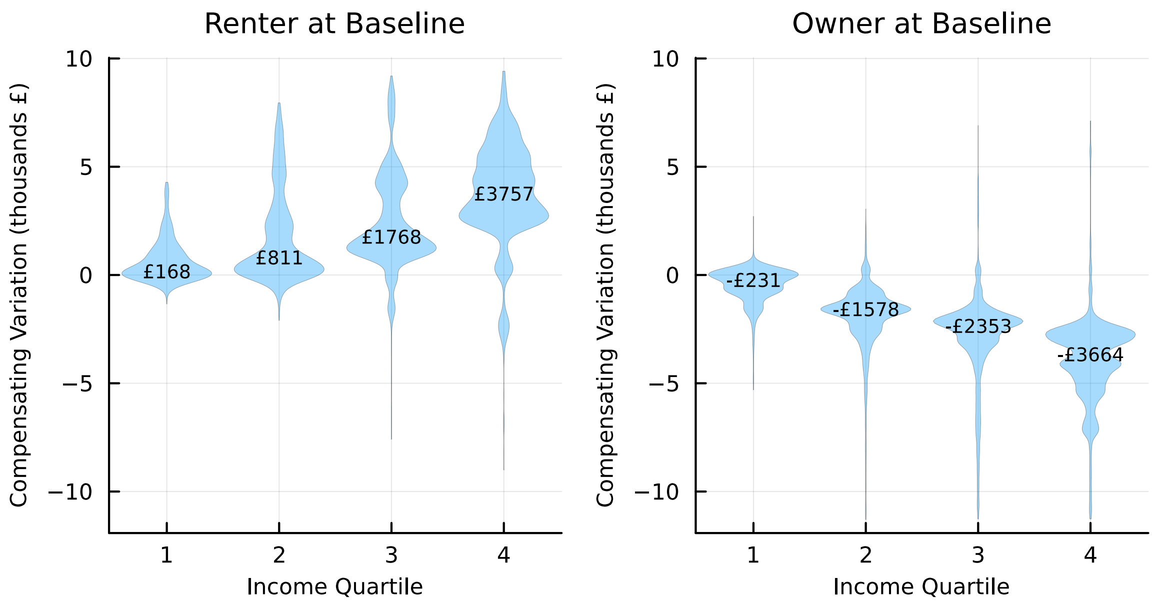

Note: This figure plots the distribution of compensating variation for households associated with increasing landlord costs in London to Birmingham levels by household income. The plot on the left presents the compensating variation for households who rented at baseline, whereas corresponds to households who owner-occupied at baseline. The median compensating variation is noted for each distribution.

effect on the equilibrium rent-to-price function of introducing a 30% tax on rental income in both London and Birmingham.

Figure 18 presents the equilibrium rent-to-price functions in both cities for the baseline and counterfactual scenarios. Two key findings emerge. First, the rent-to-price function is substantially more sensitive to the rent tax in Birmingham than in London. Second, the sensitivity of the rent-to-price function is more pronounced in the lower quality segments in Birmingham.

Figure 18: Sensitivity of Rent-to-Price to Rent Tax by City

(a) London (b) Birmingham

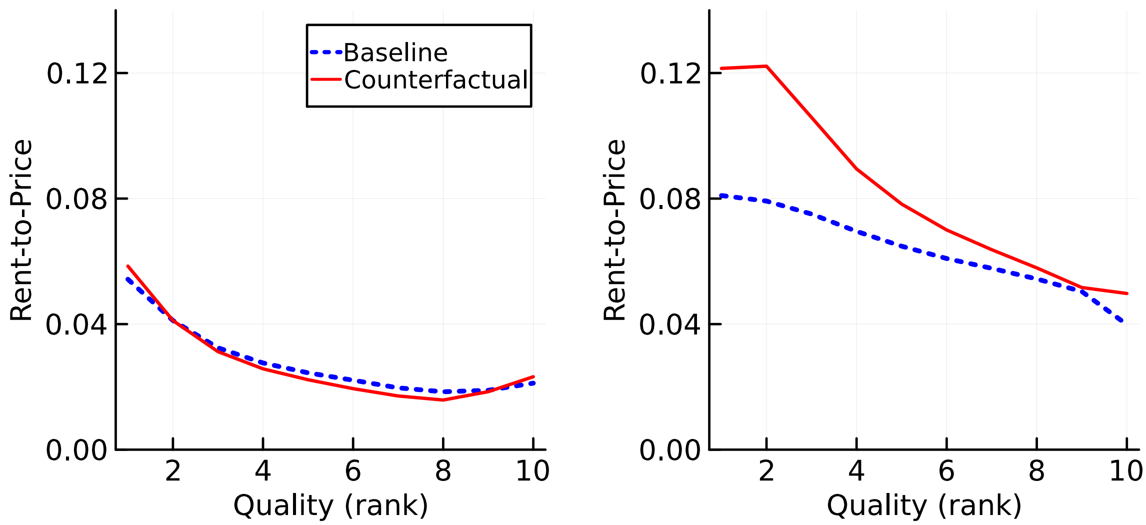

Note: This figure compares the baseline rent-to-price function with the rent-to-price function in the counterfactual in which landlords are levied an additional 30% tax on rental income.

The variation in tax sensitivity across cities and quality segments is driven by differences in landlord costs and profitability. Lower-quality segments in Birmingham exhibit greater sensitivity because landlords in these segments typically earn lower baseline profits.[32] Landlords with lower profits are more likely to pass increased tax costs onto tenants through higher rents, as they have limited capacity to absorb these additional expenses. This amplifies the tax sensitivity of the rent-to-price ratio within low-profit market segments.

The findings show that tax pass-through decreases as landlord profits and income increase. This relationship has important implications for the design of

[32] Figure 12 illustrates that landlord profits are generally lower in Birmingham than in London, with the lowest profits concentrated in Birmingham's lower-quality housing segments.

taxes for landlords. It suggests that a progressive tax, which increases with landlord income, would generate fewer distortions in the housing market than a flat tax.

2.5.4 Implications of Landlord Mobility for Counterfactual Analysis

The model does not explicitly incorporate cross-city location choices of landlords. To understand the impact of this simplification, consider a more general setting in which landlords can choose where to operate. In such a setting, they would select cities based on expected profits that reflect capital gains, rental yields, and costs. The composition of landlords would endogenously affect profits, and profits would in turn influence sorting decisions that reshape landlord composition. Equilibrium market outcomes and landlord composition would therefore be jointly determined as investors move until expected returns net of relocation costs equilibrate across locations.

The model in this paper does not include this city-choice problem. Instead, it focuses on how markets equilibrate within each city, given the composition of landlords observed in the data for that year. This simplification is mainly for tractability and estimation. This assumption does not affect the estimation of model parameters, since the yearly data on the wealth distribution of landlords in each city already reflect any relocation decisions made in previous years. The estimation solves the model conditional on the realized distribution. However, the decision to not model the city-choice problem does influence the interpretation of the counterfactual exercises. In these simulations, landlord composition is treated as fixed, which means that changes in policy or market conditions do not induce capital re-allocation across cities. The results therefore represent short-run outcomes, before landlords have time to adjust their portfolios across locations.

The importance of this missing channel is limited by the fact that most landlords are small investors who tend to own properties close to where they live. According to a survey of private landlords in the UK, 88% of landlords based in London own properties within London. These shares are also high in other regions of the UK (Scanlon and Whitehead, 2016). Similarly, evidence based on administrative data for France shows that the probability of investing in a municipality falls sharply with distance from the investor's home (Levy, 2022). Therefore, if small landlords face moving costs which make them reluctant to relocate across cities, their tendency to invest locally implies that their capital will remain largely immobile across cities.

The assumption of limited mobility is likely less appropriate for larger institutional landlords, who we do not model in this paper. These investors may face lower costs of re-allocating capital across cities and may therefore be more likely to adjust to changes in expected returns across locations. By omitting this capital re-allocation mechanism, the current model captures a short-run equilibrium in which local prices and quantities adjust while the composition of landlords remains constant.

In the longer run, investor mobility can alter the incidence and magnitude of the impact of policy. For instance, the counterfactual in Section 2.5.2 shows that the rent-to-price ratio increases in London when landlord costs are raised. While this exercise ignores re-sorting, we might expect that a relative increase in landlord costs in London would induce some re-allocation of capital toward cities such as Birmingham. This re-allocation across cities might further raise the rent-to-price ratio in London, since renters would need to offer higher rents relative to prices to incentivize landlords to continue operating in London. At the same time, the resulting increase in landlord supply can reduce the rent-to-price ratio in Birmingham.

The estimated parameters of the structural model show that baseline costs for landlords are lowest in London and highest in Birmingham. This difference is consistent with the discussion in Chapter 1 that London has greater expected capital gains and lower depreciation and risk. While this paper does not examine the underlying sources of these differences, existing work shows that spatial variation in capital gains reflects persistent differences in building supply constraints and long-run demand growth across cities (Hilber and Vermeulen, 2016; Hilber and Mense, 2025). In a general model of the housing market, a relaxation of local building supply constraints or a shift in local demand could induce landlords to re-allocate capital across cities in response to the resulting change in expected capital gains. The current model does not allow for this margin of adjustment.

2.6 Conclusion

In this paper, I analyze how landlord heterogeneity shapes housing market outcomes across cities in England. I develop and estimate a two-sided assignment model in which heterogeneous households and landlords match to differentiated housing. I find that heterogeneity in observed landlord wealth within cities and unobserved opportunity costs across cities play a key role in explaining variations in rental affordability, segmentation between rental and owner-occupied housing, and the relative supply of rental properties.

Using counterfactual experiments, I quantify the impact of landlord heterogeneity on each of these outcomes. First, I show that within-city heterogeneity in landlord wealth almost completely explains differences in the rent-to-price ratio across quality segments within a city. Second, I demonstrate that heterogeneity in landlords' unobserved opportunity costs explains 60% of the gap in quality segmentation between London and Birmingham and up to 67% of the gap in the user cost-to-rent ratio. Finally, I find that landlord heterogeneity moderates the equilibrium impact of housing policy. I show that a tax on rental income has lower pass-through in cities with wealthier landlords. These findings suggest that landlord composition is crucial in determining housing affordability and accessibility, with significant implications for housing policy.

Chapter 3

The Anatomy of a Shock to Residential Real Estate: The Role of Lenders*

3.1 Introduction

Housing and mortgage markets are repeatedly shocked by physical hazards (rising seas, extreme weather, wildfires) and by information that prompts a reassessment of the collateral risk these hazards pose (see e.g., Bernstein et al., 2019; Baldauf et al., 2020; Giglio et al., 2021). When lenders downgrade the collateral quality of homes exposed to such risks, credit flows can contract and asset prices can adjust, a dynamic long emphasized by the financial-accelerator mechanism of Bernanke et al. (1996). Yet we still lack clear evidence on which households lose access to credit, how financing gaps are bridged, and how overall housing utilisation changes when a new collateral risk emerges or is reassessed.

Because housing is the largest asset on household balance sheets and the mortgages they secure dominate bank portfolios, contractions in housing credit can widen wealth inequality, shift risk within the financial system, and alter the owner-

*This chapter is based on co-authored work with Benjamin Guin and Liam Clarke which is available as Bank of England Staff Working Paper No. 1,111. We thank Tamiko Bayliss, Alessandro Gavazza, Ben Keys, Derek Nesbitt, Felipe Netto, Arthur Taburet, Quynh-Anh Vo as well as staff at the Bank of England for their comments. Any views expressed are solely those of the authors and should not be taken to represent those of the Bank of England or as a statement of Bank of England policy. This paper should not be reported as representing the views of members of the Monetary Policy Committee or Financial Policy Committee.

ship structure of housing. Yet isolating the effects of such *risk reassessments*, often described in the climate-finance literature as transition risk[33] has proved difficult for three reasons. First, most hazards examined to date are either highly local (floods, fires) or slow-moving (sea-level rise), making it hard to disentangle collateral risk from local macro conditions or to pin down timing. Second, datasets that simultaneously track lending, transactions, and rents at scale are rare. Third, new information about risk often emerges shortly after a physical event, making it difficult to separate market responses to the event itself from responses to subsequent shifts in how risk is perceived and priced.

We address these hurdles by analysing a two-stage shock that targeted a single segment of the UK housing market, high-rise flats with combustible exterior cladding, hereafter, "cladded flats". The shock unfolded in two distinct stages. The first stage was the tragic Grenfell Tower fire of 14 June 2017, which resulted in significant loss of life and heightened concerns about the safety of high-rise cladded dwellings.[34] Fourteen months later, the Ministry of Housing issued Advice Note 14, formally identifying these flats as at risk and advising building owners to commission professional checks of their external wall systems and, where necessary, carry out costly remediation. Mortgage lenders soon treated proof of compliance as essential, with most making an EWS1 (External Wall System) safety certificate a

[33]Transition risk refers to the valuation shock that arises when new information changes how collateral is priced (e.g. Carney, 2015; Giglio et al., 2021)

[34]The authors of this paper recognize the suffering and hardship the Grenfell disaster has brought to those affected and acknowledge the significant personal and social impacts it has had, which extend beyond the scope of the empirical analysis. This paper is focused on the financial stability implications arising from shocks to residential properties. We refrain from commenting on the actions of the decision-makers involved.

de facto prerequisite for lending.[35] The fire therefore constitutes the physical stage, whereas the guidance represents the transition stage of the risk. Because exposure is determined solely by building form rather than location, the episode generates a natural experiment that isolates collateral re-assessment from broader local trends.

To quantify the effects, we combine three administrative datasets that collectively span lending, sales, and rental activity in the housing market: (i) the Product Sales Database, which covers the universe of mortgage originations; (ii) the Land Registry Price Paid dataset, which records every property transaction; and (iii) rental listings from Zoopla–WhenFresh, the second-largest online property-listings platform in the UK. Our analysis focuses on London, where high-rise cladded flats are most concentrated;[36] nationwide estimates display the same qualitative patterns, albeit with smaller magnitudes. Lacking data on building height, cladding indicators or EWS1 certificates, we conservatively classify all flats as treated, which biases estimates toward zero and thus understates the true effects.

Our empirical strategy employs a dynamic difference-in-differences design. For prices and rents we compare at-risk flats with non-flats within the same postcode and include property or postcode $\times$ rooms fixed effects. For transaction counts we aggregate to the postcode-district level and control for district fixed effects. Quarterly event-study coefficients trace each outcome from the 2017 fire through the 2018 guidance and up to just before the beginning of the pandemic, and pre-trend tests show no differential movements prior to the fire in London.

[35] Although Advice Note 14 is not statutory, most lenders swiftly conditioned new lending or refinancing on an EWS1 certificate; see https://www.insidehousing.co.uk/insight/advice-note-14-explained-what-is-it-and-why-is-it-stopping-the-sale-of-so-many-properties-63981.

[36] We define London as all postcode districts whose outward codes fall under the Royal Mail Greater London sorting area (e.g. E, EC, N, NW, SE, SW, W, WC, BR, CR, etc.).

While the assumption of parallel pre-trends holds in London, it is violated outside London, where prices appreciated more rapidly for flats relative to non-flats prior to the fire. We therefore interpret only the estimates for London as causal, whereas we treat results for regions outside London as descriptive and present them only for comparison. This focus on London is appropriate given that the Grenfell fire occurred in London and London accounts for by far the largest share of high-rise residential properties in England.[37]

We document four main results. First, mortgage originations on affected flats fell about 30 % in London (20 % nationally) once the guidance was issued. The decline was particularly pronounced for first-time buyers, who experienced a nearly 50 % drop in originations. This pattern is consistent with the financial-accelerator mechanism of Bernanke et al. (1996), whereby lenders tighten credit most for borrowers perceived to pose higher agency costs.[38] Second, the pull-back is concentrated among smaller lenders, whose originations drop by roughly 60 %, whereas the four largest banks cut lending by about 30 %. Third, the cash-buyer share for flats rises by roughly five percentage points, partly substituting for lost credit and moderating further price declines. Fourth, while prices of affected flats fall by about

[37]According to the Department for Levelling Up, Housing and Communities (DLUHC) Building Safety Data Release: August 2023, London accounts for about 61% of all high-rise residential buildings in England. The regions with the second most number of high rises is the South East which accounts for only roughly 10% of all high rises. See:
https://assets.publishing.service.gov.uk/media/6501c248702634001389b7c1/Building_Safety_Data_Release_August_2023.pdf.

[38]First-time buyers (FTBs) are costlier for lenders to screen and monitor because they lack a prior mortgage-repayment record, typically hold thinner credit files, and have smaller liquid-asset buffers than repeat buyers. These features raise agency and verification costs, so FTB credit is among the first to be rationed when perceived collateral risk increases

five percent, rents climb by three percent, providing evidence that demand to reside remains resilient even as collateral quality is downgraded.

Taken together, the evidence is consistent with a lender-driven contraction of credit supply whereby collateral became stranded, yet shelter demand, as reflected in stable or rising rents, stayed largely intact. Although we cannot rule out that our results partly reflect concurrent reductions in demand to own flats among mortgage-dependent buyers, the timing and sequence of events following the shocks are difficult to reconcile with a purely demand-based interpretation, as discussed in Section 3.5. In the absence of data that would allow us to quantify the relative importance of supply and demand channels, we interpret the evidence as indicating that reductions in lending likely played an important role in shaping the housing market response, while acknowledging that some portion of the estimated effects may reflect demand adjustments.

The documented patterns suggest a risk–equity trade-off that, to our knowledge, has received little empirical attention in the transition-risk and climate-finance literature. When a safety shock or disaster casts doubt on the viability of a class of properties as collateral, banks can off-load risk from their balance sheets by curtailing new lending, while cash-rich buyers can capture a larger share of transactions at discounted prices. Credit-constrained households, especially younger, would-be first-time buyers, may continue to reside in the same dwellings as tenants, so the shock reshuffles ownership rather than use, delaying their wealth accumulation while leaving aggregate housing service utilization unchanged.

In our setting the distributional consequences are likely amplified because flats constitute the entry point of the UK home-ownership ladder. Postponing acquisition of this first rung can cascade through subsequent moves and widen lifetime wealth gaps, consistent with the housing ladder dynamics emphasized by Ortalo-Magné and Rady (2006). A comparable pattern emerged after the 2008 U.S. foreclosure wave, when cash-rich investors absorbed distressed stock, stabilised prices, and deepened the long-run drop in home-ownership. Our evidence shows that a similar dynamic can emerge when a salient shock raises perceived physical or environmental risk for a class of properties, prompting lenders to retrench even as the demand to reside in dwellings remains firm.

Recent studies document that capital continues to reach hazard-exposed housing and risk is not fully capitalized into prices (e.g. Bernstein et al., 2019; Baldauf et al., 2020). We show that cash buyers can help sustain such capital flows. When mortgage credit dries up, cash buyers acquire the units and recoup the extra risk through higher rents, keeping occupancy high even as ownership shifts away from credit-constrained households. Recognising that lenders and cash buyers react very differently to the same risk signal is therefore crucial for predicting how utilisation and the distribution of housing wealth will evolve after future safety or climate disclosures.

3.1.1 Related Literature.

Our paper contributes to the literature on the role of credit in propagating shocks to the economy (Kiyotaki and Moore, 1997; Brunnermeier, 2009; Bernanke et al., 1996). Consistent with the concept of credit serving as a “financial accelerator" (Bernanke et al., 1996), we observe a decrease in collateralized mortgage lending to households following an adverse shock to residential property values. In line with this notion, our paper complements existing evidence on the relationship between property values and lending to firms (e.g., Chaney et al. (2012); Gupta et al. (2021)). Recent research indicates that firms seek new debt following an increase in the value of their real estate, but they tend to opt for unsecured borrowing rather than secured borrowing (Campello et al., 2022). Consistent with the notion of reduced risk appetite,[39] our findings indicate that the decline in mortgage lending is most significant for informationally-opaque borrowers. Specifically we find evidence of a decline in mortgage lending to First-time-buyers (FTB) that are more informationally opaque given their shorter credit history. This underscores the role of information asymmetries between lenders and borrowers (Sufi, 2007) in constraining lending (Stiglitz and Weiss, 1981). In this context, our findings on residential mortgages complement research on corporate lending, which also demonstrates evidence of credit rationing among younger firms (Kirschenmann, 2016; Ferri and Murro, 2015). Overall, our study provides valuable evidence on residential mortgage lending which comes with implications for household finance, expanding our understanding beyond the scope of existing literature focused primarily on lending to firms.

The established view of real estate collateral’s role in economic fluctuations has faced recent challenges from empirical studies (e.g., Mian and Sufi (2011)).[40]

[39] In line with the concept of a “flight for safety" (Bernanke et al., 1996).

[40] They argue that the primary transmission channel of financial shocks to the real economy during the Great Recession was the significant reduction in aggregate demand driven by declines in household net worth.

We contribute to this literature by introducing two additional perspectives. Firstly, we highlight the crucial role of cash-buyers and lender heterogeneity in mitigating the effect of adverse shocks to real estate. Previous research has suggested that heterogeneity in risk beliefs can explain why drops in property values are less severe than expected (Baldauf et al., 2020; Bakkensen and Barrage, 2021). We complement this literature by providing evidence that the heterogeneity in financing sources for property purchases also plays a relevant role in understanding shocks to property prices. Specifically, our findings suggest that the proportion of cash-buyers in the property market is a significant factor influencing the impact of shocks on property values.

Our paper also contributes to the literature on how lenders respond to such shocks to real estate which acts as collateral for mortgage lending.[41] This literature documents heterogeneous responses with respect to banks' levels of capitalization (Schüwer et al., 2018) and diversification (Chavaz, 2016). We contribute to this literature by documenting nuanced differences across bank size where small lenders react more strongly after the shock while larger lenders' lending decreases less strongly (potentially because they are more diversified). We also complement to it by documenting a lagged effect on mortgage lending which dropped only following public

[41]The existing literature has examined responses to tightening credit limits and interest rates (Cerqueiro et al., 2016; Nguyen et al., 2022), securitization decisions (Ouazad and Kahn, 2021) and property valuations (Garbarino and Guin, 2021).

guidance is issued but not immediately after the physical shock.[42]

3.2 Overview of the shock and data

3.2.1 Timeline of events

In this paper, we study market dynamics around a specific shock to properties in the UK, which encompasses two events that occurred consecutively. They break our sample up into three periods (Figure 19).

Figure 19: Timeline of Events

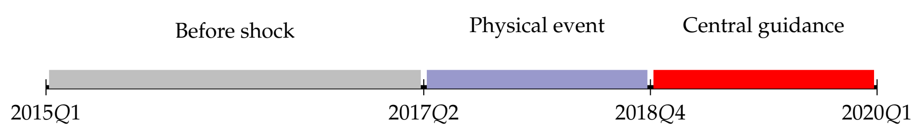

The first event was a severe fire that affected a specific building located in London which resulted in the tragic loss of life.[43] This disaster raised concerns

[42]Our paper also links to the literature studying dynamic responses to shocks to real estate both in terms of prices and transaction volumes. It shows that there is a lead-lag structure which means property transaction volumes fall first and then prices only fall later in the presence of a slow-moving shock from sea level rises (Keys and Mulder, 2020) which can be rationalized by optimistic homeowners continue listing properties at high price and only gradually update their asking prices which explains the sluggish drop in prices (DeFusco et al., 2022). By contrast, our results suggest in the presence of an unexpected and salient shock both transaction volumes and prices fall immediately. For the mortgage market, our results suggest in fact a reversed lead-lag relationship, i.e. valuations fell immediately whereas volumes responded more than a year later, suggesting that mortgage lenders update their valuations immediately.

[43]The Grenfell Tower Inquiry was created to examine the circumstances leading up to and surrounding the fire at Grenfell Tower on the night of 14 June 2017. More information can be found at `https://www.grenfelltowerinquiry.org.uk/`

about the safety of flats in other high-rise buildings because of cladding which was identified to be highly combustible. It also had further, significant personal and social impacts on peoples' lives, which extend beyond the scope of our academic analysis. The second event occured in December 2018 when the government released central guidance, the so-called Advice Note 14. From the perspective of our analyses, it served two purposes. First, it defined the subset of flats that were at risk and those that were not, specifically those high-rise flats with exterior wall cladding. Second, it suggested remedies for these affected properites. For building owners, it mentioned the measures they should take to ensure their buildings were safe. Specifically, it guided owners of flats to check whether their property was constructed using unsafe materials and to remove such material if found, which comes with costs.[44] For the purpose of our empirical analyses, we restrict our sample to start in 2015 Q1.[45] We choose to end the sample period in 2020 Q1 because of the first COVID lockdown in the UK that ocurred at the time which may conflate our analyses.

3.2.2 Data

In our analysis, we employ transaction-level data from the Product Sales Database (PSD). It records the universe of residential mortgage originations and refinancings in the UK during our period of interest and it is well-established in the literature (e.g., Benetton (2021); Benetton et al. (2021); Peydró et al. (2023); Arnould et al. (2020)). In terms of geographical information, we observe the six-digit postcode of

[44]Such costs include (higher) insurance costs for affected properties as well as costs of maintenance and renovations.

[45]Choosing 2015 Q1 ensures that the time before the event roughly equals the time after the event.

each mortgage transaction.[46] In our sample of mortgage originations in England & Wales between 2015 and 2020, about 13% of all mortgage transactions are flats (remaining ones are non-flats). We observe on average 19 transactions of non-flats and 3 transactions of flats per postcode district and month before the shock. The average price paid for a mortgage-financed flat is £283,565 and the average price of a mortgage-financed non-flat is £296,667. We use these data to quantify the number of property transactions and average prices by property type, geographic area and time. We employ postcode district as the geographic area and the year-month level as our time dimension. In terms of property type, our data allows us to differentiate between flats and non-flats. However, we are neither able to observe whether flats are in high-rise buildings nor can we observe their wall construction such as cladding.

Whilst the PSD provides rich data to study trends in property prices and the number of mortgages originated, it does not contain information on those property transaction that are cash-financed. We therefore combine PSD with the Land Registry Price Paid Dataset (PPD) to separately analyse the dynamics in the cash- and mortgage-financed segments of the housing market. The latter records the universe of all residential property transactions in England and Wales and it is also well-established in the existing literature (e.g., Bracke et al. (2018)). For each transaction, we observe the property type, the exact transaction date as well as the price paid for the property. In our sample of all transactions in England & Wales between 2015 and 2020, about 19% of all property transactions are flats (remaining ones are non-flats). We observe on average 26 transactions of non-flats and 6 transactions of flats per postcode district and month before the shock. The average price paid for a property

[46] A six-digit postcode can consist of around 10 properties.

is £289,947 in the sample period. The average price for a flat is £293,332 and the average price of a non-flat is £289,163 in the Land Registry.

Combining PPD with PSD allows us to back out the implied number of all cash purchases for flats and non-flats at the postcode district $\times$ year-month level by subtracting the number of mortgage transactions in PSD from the number of property transactions in PPD.[47] Between 2015 and 2020, the share of cash transactions among all properties was approximately 37%. This number compares well with aggregate statistics[48] and the share reported by Bracke and Tenreyro (2021) for the UK.

We complement these two datasets with data on rental listings from Zoopla which are processed by WhenFresh Ltd.[49] It is the second most popular property portal in the United Kingdom in terms of traffic and covers 70% of the whole-of-household privately rented housing stock in the UK (Bracke, 2021). The dataset contains information on the address of properties, listing sale/rental prices, and property attributes (such as property type and number of bedrooms). This dataset includes listings of rentals including the asking rent but we do not observe whether a specific property was actually rented. In the Zoopla data between 2015 and 2020 for London, the median weekly rent was £375 for all properties, £369 for flats and £398 for non-flats (see Table C.1 in the Appendix for summary statistics).

[47] PPD classifies transactions as 'Additional Price Paid' if they were "transfers under a power of sale/repossessions, buy-to-lets (where they can be identified by a mortgage), transfers to non-private individuals and sales where the property type is classed as 'Other'." Importantly, a Buy-to-Let transaction is not classified as Additional Price Paid if it was financed entirely with cash. The set of properties not tagged as Additional Price Paid therefore include all standard residential transactions (financed either with cash or mortgages) as well as Buy-to-Let transactions financed only with cash. The number of all cash purchases (including Buy-to-Let) can thus be backed out by taking the difference in the number of PPD non-Additional Price Paid transactions and all PSD transactions.

[48] www.gov.uk/government/news/uk-house-price-index-new-data-reveals-number-of-cash-buyers, retrieved on 24 September 2023.

[49] *WhenFresh Ltd* is a major provider of UK residential property data (`https://whenfresh.com`).

3.3 Research Design

For our empirical analyses, we employ a dynamic Difference-in-Difference (DiD) research design. It allows us to estimate the effects of the shock on outcomes in the property, mortgage and rental markets. Specifically, we examine how outcomes of properties affected by the shock compare to outcomes of unaffected properties before and after the shock. The long time dimension and high frequency of our data allows us to examine the dynamic effects of the shock. Specifically, we can study when exactly the effects of the shock materialized and whether they were temporary or persistent.

Employing this dynamic DiD research design is appealing given the nature of the shock in question and the data at our disposal. First, the shock was unexpected and was unprecedented.[50] This mitigates concerns that market participants' anticipation might explain some of the effects on outcomes in the property and mortgage markets.[51] Second, our setting and data allow us to identify a set of properties affected by the shock, our treatment group, which we can compare to a set of unaffected properties which form our control group.[52] As the treatment is assigned by the type and not by the spatial location of the property, concerns of changing local

[50]We argue that the shock that we study was unexpected for two reasons. First, prior to it, there was not much coverage of the risks associated with cladding issues which were responsible for it. Second, there was not much coverage of the fact that these cladding issues were particularly problematic for flats as opposed to other types of buildings.

[51]By contrast, in other settings news of expected changes to policy or risk is available well in advance of the actual shock. It is difficult to identify the effects of the actual event as individuals may respond in advance in anticipation of it. This issue is particularly relevant for climate change since many climate-related events such as changes to sea levels, temperature, and rainfall can be forecasted by households.

[52]This is because the shock only increased risk perceptions and costs associated with living in high-rise flats while leaving non-flats unaffected.

economic conditions explaining some of the estimated effects are mitigated.[53] Last, the dynamic DiD allows us to estimate the incremental effects of the two subsequent events, the physical shock and central guidance, of the shock that we study.

To implement the dynamic DiD, we estimate two types of regressions. The first set of regressions allows us to study outcomes at the transaction level, for example changes in the transaction prices of affected properties. The second set of specifications allows us to study changes to the number of transactions at the postcode district $\times$ time $\times$ property-type level. We estimate our regressions for transactions both in London as well as outside London, as dynamics and the share of affected properties across these markets might differ. In all regressions, the reference period that we compare outcomes to is 2017Q2, the quarter when the physical event occurred.

The DiD regressions at the transaction level take the following generic form:

$$\ln y_i = \beta_0 + \sum_{\tau=2015Q1}^{2020Q1} \alpha_\tau \, \mathbb{1}(i \in \text{flat}, t(i) = \tau) + \gamma_{\text{flat}} \mathbb{1}(i \in \text{flat}) + \gamma_{t(i)} + \rho(i) + \varepsilon_i \quad (34)$$

where i indexes each transaction; y_i is the outcome variable of interest; $\gamma_{t(i)}$ is an indicator for quarter-year $t(i)$ for transaction i; α_τ is the difference-in-difference coefficient for quarter-year τ; and γ_{flat} is the fixed effect for flats. Depending on the exercise, $\rho(i)$ is either a fixed effect at the property level or at the postcode $\times$ number of rooms level for transaction i. Of interest are the estimates of α_τ which measure

[53]By contrast, for instance, sea-level rise impacts all coastal properties. Similarly, floods, droughts, heat waves, and forest fires tend to impact all properties within a large catchment area. Failure to account for such spatially granular time-trends can bias estimates of effects if the treated and control units are located in different locations and these locations face different time-trends.

the evolution of outcomes of affected properties relative to unaffected properties, e.g. prices of flats relative to non-flats.

To study the evolution of the number of transactions, we estimate Poisson regressions with the dependent variable defined as the count of transactions of flats or non-flats at the postcode district $\times$ month-year level. The regression specification takes the following form:

$$\ln \mu_{ift} = \beta_0 + \sum_{\tau=2015Q1}^{2020Q1} \alpha_\tau \times \mathbb{1}(f = \text{flat}, t \in \tau) + \gamma_\tau + \gamma_f + \rho(i) \tag{35}$$

where i indexes each postcode district (our measure of the local area); t indexes each year-month (the time dimension); f indexes flats and non-flats (the property type); μ_{ift} is the expected count of the relevant transactions of type-f properties during month-year t; γ_τ is an indicator for quarter-year τ; γ_f is the fixed effect for properties of type-f; $\rho(i)$ is the postcode district fixed effect; and α_τ is the difference-in-difference coefficient for quarter-year τ.

3.4 Results

3.4.1 Effect on property prices

This section studies how the shock which is composed of two subsequent events, the physical event and the central guidance, affected transaction prices for flats. We find that prices (and therefore the value of mortgage collateral) for flats fell significantly and persistently relative to non-flats following the physical event.

3.4.1.1 Evolution of property prices over time

We first examine the historical price dynamics across the two types of properties, flats and non-flats. Specifically, we study the evolution of mean purchase prices in log units for flats and non-flats that were purchased using a mortgage. To account for heterogeneity in price dynamics across regions, we illustrate the evolution of prices separately for transactions within London and outside of it. The reported, estimated means adjust for time trends as well as compositional differences due to unobserved and observed characteristics at the postcode $\times$ number of rooms level which might influence the price of the transacted properties. Specifically we estimate four different regressions, for flat and non-flat transactions both within and outside London.

$$\ln p_i = \beta_0 + \gamma_{t(i)} + \rho(i) + \varepsilon_i \tag{36}$$

where i indexes each transaction, $\gamma_{t(i)}$ is an indicator for quarter-year $t(i)$ for transaction i, and $\rho(i)$ is the postcode $\times$ number of rooms fixed effect for i. All regressions presented in this paper exclude the quarter-year indicator for 2017Q2, the quarter when the physical event occurred. Therefore all quarter-year parameters γ_t represent deviations from the mean in 2017Q2. We report robust standard errors unless stated otherwise.

The granular postcode $\times$ number of rooms fixed effects allow us to control for any time-invariant characteristics which generate price differences across postcodes, property sizes, and the interaction of postcode and property size. These fixed effects isolate any differences in prices that may occur due to differences along these dimensions. The evolution of the estimated γ_t thus captures quarterly changes in prices due to factors other than differences in the composition of transacted properties along these dimensions. Figure 20 plots estimates of γ_t and their confidence intervals using data from the Product Sales Data 001.

Figure 20: Log Purchase Price for Mortgage-Financed Property Transactions

(a) London (b) Outside London

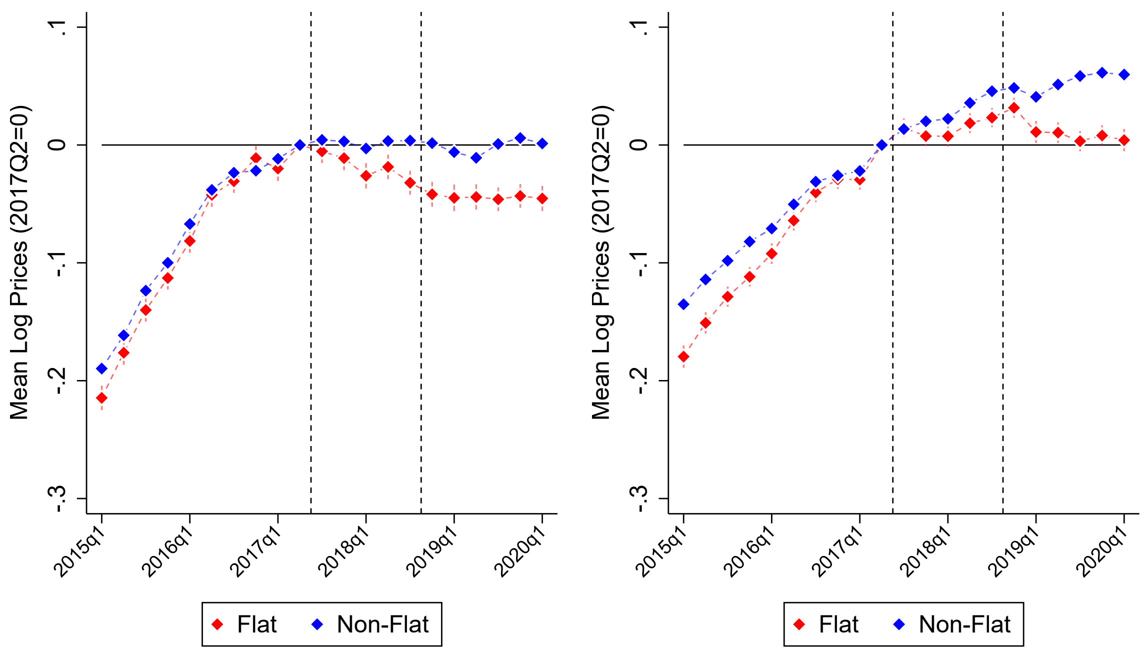

Note: This figure presents estimates from two different regressions. Each figure plots the quarterly mean log purchase price for transactions of flats and non-flats in London and outside London respectively along with the 95% confidence interval estimated from a normal distribution. The estimates adjust for fixed effects at the postcode × number of rooms level. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

The estimates in Figure 20 show an abrupt change in the evolution of flat prices relative to non-flat prices beginning in 2017Q2, the quarter of the physical event. Whereas prior to 2017Q2 in London the percentage difference in prices between flats and non-flats was weakly decreasing, this percentage difference abruptly started increasing after 2017Q2. The relative decline in flat prices continued for 6 quarters after which the percentage difference between non-flats and flats stayed approximately constant at around 6%. Figure 20 shows that transactions outside London also exhibit an abrupt change in the relative evolution of flat and non-flat prices in 2017Q2. Whereas prior to 2017Q2 the percentage difference between flat and non-flat prices was falling due to significantly greater growth in flat prices, this pattern was reversed in 2017Q2 when the growth in non-flat prices suddenly began to outpace the growth in flat prices. As with London, by 2020Q1, the percentage difference in prices between non-flats and flats had reached approximately 6% for transactions outside London.

3.4.1.2 Effect on property prices: Difference-in-Difference analyses

Our previous descriptive analyses of the evolution of transaction prices suggested that the price growth for flats outpaced non-flats prior to 2017Q2. By contrast, after 2017Q2 this pattern reversed abruptly. In this subsection, we formally test these differences in the relative change of transaction prices. To that end, we conduct a dynamic difference-in-differences analysis. We estimate the following specification

Figure 21: Flat vs. Non-Flat Difference in Log Purchase Prices
for Mortgage-Financed Property Transactions

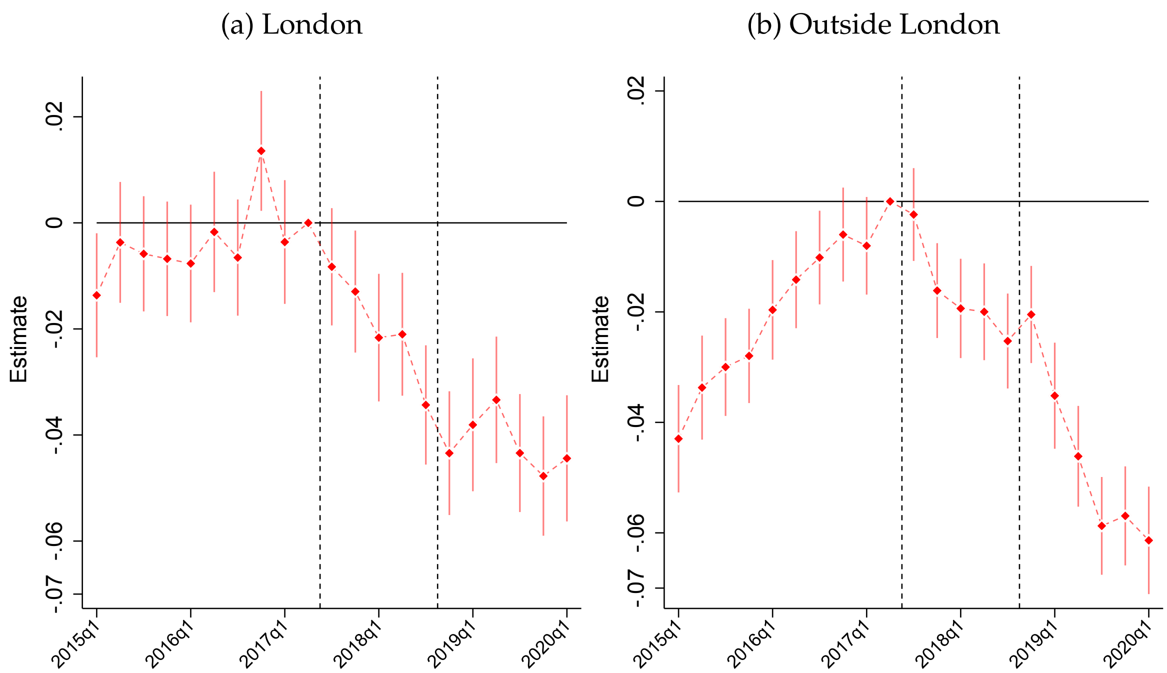

Note: This figure presents difference-in-difference estimates from two different regressions. Each figure plots the quarterly difference in mean log purchase price between mortgage-financed flat and non-flat transactions in London and outside London respectively along with the 95% confidence interval estimated from a normal distribution. The estimates adjust for fixed effects at the postcode × number of rooms level. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

separately for transactions within and outside London:

$$\ln p_i = \beta_0 + \sum_{\tau=2015Q1}^{2020Q1} \alpha_\tau \, \mathbb{1}(i \in \text{flat}, t(i) = \tau) + \gamma_{\text{flat}} \mathbb{1}(i \in \text{flat}) + \gamma_{t(i)} + \rho(i) + \varepsilon_i \quad (37)$$

where i indexes each transaction; $\gamma_{t(i)}$ is an indicator for quarter-year $t(i)$ for transaction i; $\rho(i)$ is the postcode × number of rooms fixed effect for i; α_τ is the difference-in-difference coefficient for quarter-year τ; and γ_{flat} is the fixed effect for flats. Of our interest are the estimates of α_τ which gauge how much prices of flats have evolved compared to non-flats relative to the 2017Q2, the quarter of the physical event.

Figure 21 plots these dynamic difference-in-difference estimates. They provide empirical evidence that flat prices reversed abruptly relative to non-flats in 2017Q2. In London, the point estimates for the difference in mean log prices are increasing before 2017Q2, although they are not statistically significantly different than zero at the 5% level. We conclude that there is no statistical evidence for the existence of a pre-trend in the difference in log prices between flats and non-flats. However, the difference in log prices between flats and non-flats begins to abruptly fall after 2017Q2, such that within two quarters (2017Q4) the difference is statistically significantly negative and different from zero. By 2020Q1 the difference in mean log prices has fallen statistically significantly to approximately -0.05 log points which represents a 5% drop in the prices of flats relative to 2017Q2 and the quarters prior.

Outside London, there is evidence for a reversal of the relative evolution of prices between flats and non-flats. Unlike in London, in the time period before the physical event, we see a strong and statistically significant increase of 4% in prices of flats relative to non-flats. Yet, this pattern completely reverses in 2017Q2 such that flat prices drop statistically significantly by $\sim$2% within two quarters (2017Q4) and by $\sim$6% by 2020Q1. Due to the presence of differential pre-trends, we do not interpret the evidence for outside London as representing the casual effect of the fire.

Overall, this evidence presented so far provides strong evidence that the physical event reduced prices for flats purchased with mortgages by approximately

5% in the span of two years in London. On an annualized basis, the physical event reduced flat prices by approximately 2.5% per year.

3.4.2 The role of lending

Consistent with our priors, our previous analyses suggest that there is an economically significant and persistent negative effect of the shock on the prices of affected properties relative to unaffected ones. In this section, we examine the dynamics of mortgage lending which serves to finance property purchases. Reduced market values of properties reduces the value of collateral where there is an outstanding mortgage. Lenders might respond to uncertanity of property values by reducing mortgage originations against affected properties.

3.4.2.1 Contraction in mortgage lending against affected properties

First, we examine whether mortgage lending against affected properties changed in response to the shock. To that end, we study the relative change in the number of mortgage transactions when affected properties serve as collateral compared to non-flats. As our dependent variable is a count variable, we employ a Poisson regression model. Specifically, we estimate the following dynamic difference-in-differences specification:

$$\ln \mathbb{E}[Q_{ift} \mid \cdot] = \beta_0 + \sum_{\tau=2015Q1}^{2020Q1} \alpha_\tau \times \mathbb{1}(f = \text{flat}, t \in \tau) + \gamma_\tau + \gamma_f + \rho(i) \qquad (38)$$

where i indexes each postcode district; t indexes each year-month; f indexes flats and non-flats; Q_{ift} is the count of originations of type-f properties during month-year t; γ_τ is an indicator for quarter-year τ; γ_f is the fixed effect for properties of type-f; $\rho(i)$ is the postcode district fixed effect; and α_τ is the difference-in-difference coefficient for quarter-year τ. We present the difference-in-difference estimates of α_τ. We estimate this regression for two different subpopulations, both within London as well as outside London (Figure 22).[54]

The difference-in-differences estimates presented in Figure 22 document a sharp drop in mortgage originations against flats in 2018Q4, the quarter when the central guidance was published across both regions. However, the drop is significantly more pronounced in London where the number of mortgages against flats fell sharply and persistently by 30-40%. In comparison, flat originations dropped by around 10-20% outside London. The fact that originations dropped by an additional $\sim$20 percentage points in London is consistent with ex-ante expectations since the central guidance targeted high rise properties which are proportionally more common among flats in London.[55]

By contrast, estimates of the 19 quarters preceding 2018Q4 do not suggest anticipation effects because estimates before 2018Q4 are close to zero (with the only statistically significant departure from zero being due to yearly seasonal trends which drive up the estimate predictably in the first quarter of each year). The exact coincidence of the central guidance with this significant and sustained departure from the pre-2018Q4 trend provides strong evidence that this drop in flat originations was

[54] We also document results for England & Wales, as shown Figure C.1 in the Appendix.

[55] This evidence further strengthens our interpretation that the post-2018Q4 drop was indeed caused by the central guidance as it rules out any competing explanations for the drop in 2018Q4 that cannot also simultaneously explain the pattern of heterogeneity documented in Figure 22.

Figure 22: Flat vs. non-Flat Difference in Log Number of Originations by Region

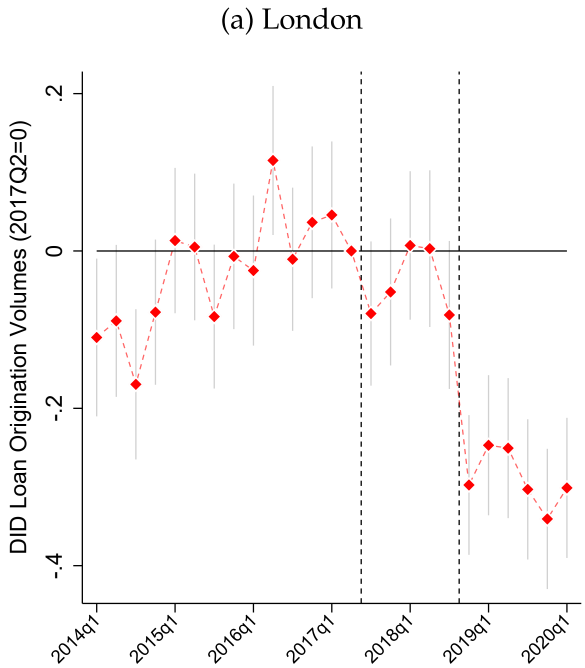

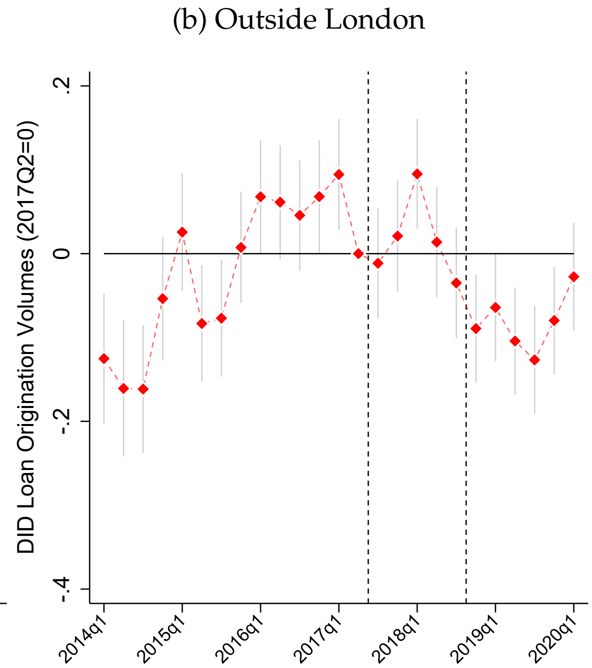

Note: This figure presents difference-in-difference estimates from two different Poisson regressions. Each figure plots the quarterly difference in mean log number of mortgages originated for flats and non-flats within and outside London along with the 95% confidence interval estimated from a normal distribution. The estimates adjust for postcode district fixed effects. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

caused by the central guidance. In contrast, the data shows no significant evidence that the physical event had an impact on mortgage originations.

3.4.2.2 Less credit to First-time buyers (FTB) that are information opaque?

In this section, we explore the heterogeneity of the shock on mortgage lending to First-time buyers (FTB) given their shorter credit histories.[56] To that end, we estimate

[56]We compare them to home movers that have a longer credit history.

Equation 38 on the subsamples of First-time buyers (FTB), which we compare to Home Movers. The estimates in Figure 23 shows the quarterly difference in mean log number of mortgages originated for flats and non-flats in England & Wales by these two types of borrowers, Home movers vs. First-time buyers (FTB). The results suggest that flat originations dropped by close to 50% for First-time buyers (FTB). By contrast, the drop of mortgage lending to home movers amounted to $\sim$30%. Overall, this evidence suggests that mortgage lending to First-time buyers (FTB) decreased more. This contractionary response in mortgage lending can be indicative of delaying entry into the property ladder for young home buyers that are most credit-constrained and reliant on mortgage finance.

We acknowledge that first-time buyers usually have higher loan-to-income and loan-to-value ratios as they tend to borrow a higher proportion of properties' costs. Given this positive correlation, the difference in originations documented in Figure 23 might simply reflect differences in leverage between the two types of borrowers. To address this concern, Appendix Figures C.2 and C.3 explicitly check whether the drop in origination volumes for affected properties are correlated with the LTV and LTI ratios. We find no statistically significant evidence for differential effects with respect to leverage.

Our results point to credit-dependent buyers—especially first-time buyers—as those most affected by the shock. The properties in question were mostly flats, which typically serve as the entry point to home-ownership in the UK. Disruptions at

Figure 23: Flat vs. Non-Flat Difference in Log Number of Mortgage Originations in London, by Type of Borrower

(a) First-time Buyers (FTB)

(b) Home Movers

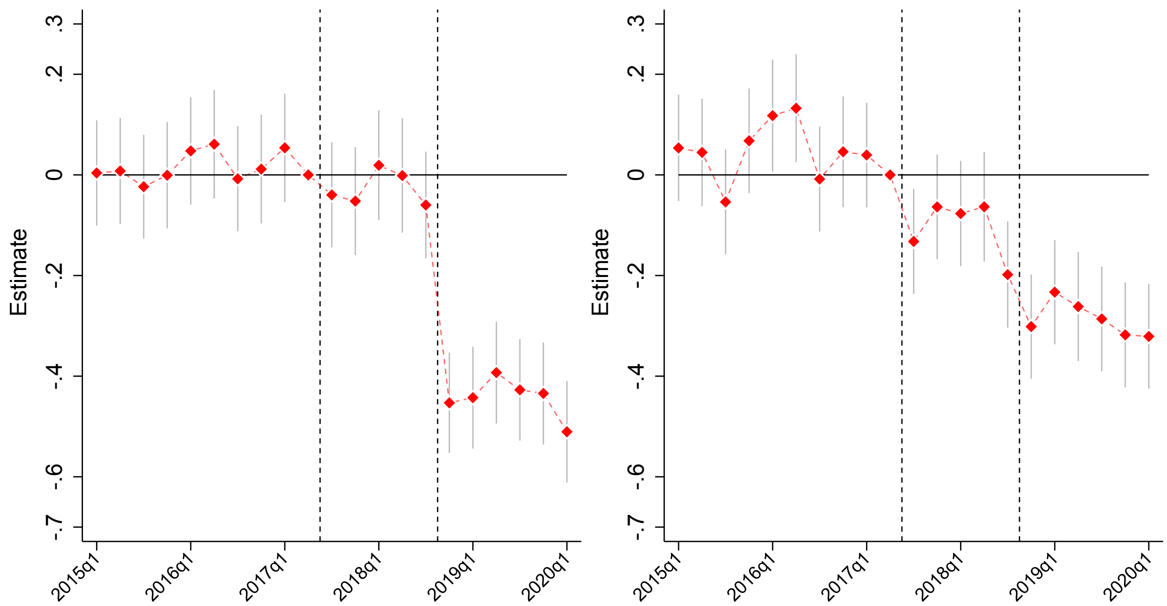

Note: This figure presents difference-in-difference estimates from two different Poisson regressions. Each figure plots the quarterly difference in mean log number of mortgages originated for flats and non-flats in London by different types of borrowers along with the 95% confidence interval estimated from a normal distribution. The estimates adjust for postcode district fixed effects. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

this stage risk delaying or blocking access to the housing ladder, with potential consequences for long-term housing trajectories and wealth accumulation.[57]

[57]The shock may also have made it harder for existing owners to trade up, by limiting their ability to sell affected properties.

3.4.3 Mitigants to the effects of the shock

3.4.3.1 Heterogeneous response by lender size

Figure 24: Flat vs. Non-Flat Difference in Log Number of Mortgages Originated by Lenders of Different Sizes in London

(a) Small Lenders

(b) Large Lenders (Top 4)

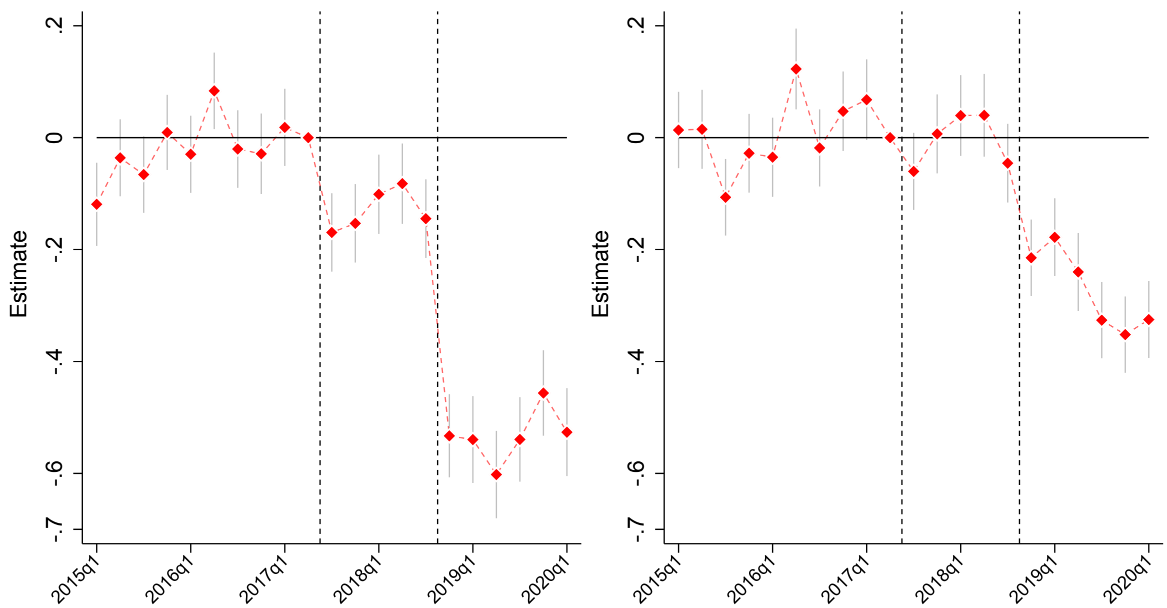

Note: This figure presents difference-in-difference estimates from two different Poisson regressions. Each figure plots the quarterly difference in mean log number of mortgages originated for flats and non-flats in London by lenders of different sizes along with the 95% confidence interval estimated from a normal distribution. The estimates adjust for postcode district and lender fixed effects. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

Our previous results show that mortgage lending dropped sharply following the central guidance. Yet, these results do not shed light on whether this drop in originations was common across mortgage lenders. There are reasons to expect that large lenders react differently to the shock than small lenders. Their mortgage portfolios might be better diversified. Moreover, they might be in a better position to absorb shocks to collateral values. For these reasons, their reaction to the shock might be less sensitive.

In this section, we examine heterogeneity in the number of mortgage originations against affected properties by lender size. In England & Wales, the top four lenders take up the largest fraction of the residential mortgage market. Specifically, the mortgage market for flats is highly concentrated among these lenders who account for approximately 90% of mortgage originations against flats and 50% of mortgage originations against non-flats. The remaining lenders account for the residual share of mortgage lending.

To capture potential differences by lender size, we estimate two separate regressions, one for the top four lenders and another one for all other lenders. To study the heterogeneity in mortgage originations by lender size, we estimate a slightly augmented and more granular version of the Poisson regression in Equation 38. The regressions take the following form:

$$\ln \mathbb{E}[Q_{iflt} \mid \cdot] = \beta_0 + \sum_{\tau=2015Q1}^{2020Q1} \alpha_\tau \times \mathbb{1}(f = \text{flat}, t \in \tau) + \gamma_\tau + \gamma_f + \gamma_l + \rho(i) \tag{39}$$

where i indexes each postcode district; t indexes each year-month; f indexes flats and non-flats; l indexes each lender; Q_{iflt} is the count of originations of type-f properties originated by lender-l during month-year t; γ_τ is an indicator for quarter-year τ; γ_f is the fixed effect for properties of type-f; γ_l is a lender fixed effect; $\rho(i)$ is the postcode district fixed effect; and α_τ is the difference-in-difference coefficient for quarter-year τ. Figure 24 plots estimates of α_τ and its confidence intervals for London.

The estimates in Figure 24 reveal that following the central guidance, small lenders reduced flat originations by nearly ∼60% while the top four lenders only reduced originations up to ∼40%. Overall, this difference of ∼20 percentage points suggests that large lenders' mortgage lending is less sensitive to the shock than lending by small lenders. As credit availability is relevant for buyers taking out mortgage debt to purchase properties, we take these results as evidence that lender size being a relevant dimension in the propagation of the shock to property values.

3.4.3.2 The role of cash-buyers

Our previous results suggest that mortgage lending dropped significantly in response to the shock. Yet, debt-financed purchases make up only around two thirds of the property market. In this section, we examine whether the drop in mortgage lending against affected properties was partially or fully offset by cash-buyers that do not take out mortgage lending to finance their property purchase.

To that end, we construct a new measure for cash purchases as such data is not publicly available. Specifically, we combine data on property transactions with those on mortgage transactions. First, we use data from the Land Registry Price Paid data (PPD) to construct a count of all flat and non-flat transactions respectively that occurred in each month-year in each postcode district. Second, we use mortgage data from the Product Sales Data (PSD) to construct analogous counts for mortgage-financed transactions.[58] We then take the difference between the first set of counts measuring all transactions and the second set of counts measuring only mortgage-financed transactions to yield a new set of counts which measure the implied number

[58] These counts of mortgage transactions are the same counts we use when estimating Equation 38.

Figure 25: Flat vs. Non-Flat Difference in Cash Share of Sales in London

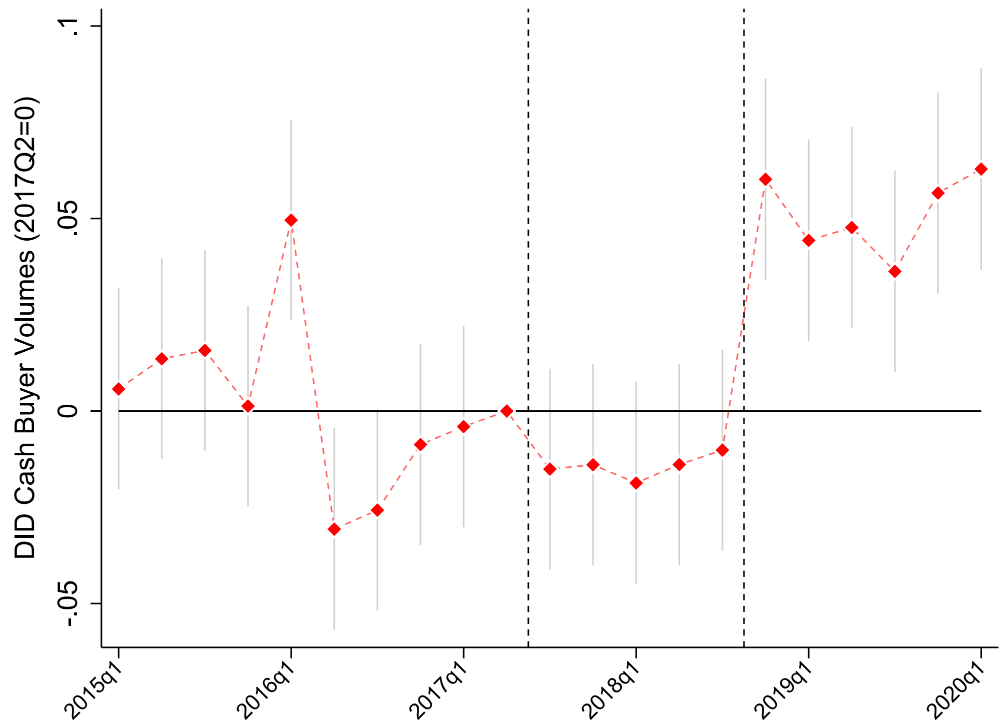

Note: This figure presents difference-in-difference estimates. The figure plots the quarterly difference in mean cash share of sales for flats and non-flats in London The estimates adjust for postcode district fixed effects. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

of cash transactions at the postcode district×month-year level for flats and non-flats respectively. The count of these cash transactions allows us to calculate the share of cash transactions relative to all transactions per postcode district.

We employ this share variable as the dependent variable in a regression estimated using the subsample of properties located in London.[59] In particular, we estimate a version of Equation 38 where the dependent variable is defined to measure

[59] We restrict the sample geographically given the size of our final data set.

the share of purchases of properties of type-f (flats or non-flats) in postcode district i during month-year t that is financed with cash. These estimates are presented in Figure 25. It shows that the share of cash-transactions against affected properties increased by approximately 5% following the central guidance relative to unaffected ones.

The results show that there is a significant increase in the share of flat purchases financed solely with cash. This result suggests that the contraction in the financing of flat purchases was offset by' financing from another source, namely, by cash-buyers. Combining these results with those on the impact on first-time buyers in section 3.4.2.2, our analyses suggest that there is reshuffling in the composition of the residential property market. Specifically, we observe that the decrease in first time buyers (FTB) coincides with increase in cash buyers. This suggests that some FTB might have been crowded out of the property market.

3.4.3.3 Rental market: Increases in rents

Our previous results illustrated the propagation of the shock on prices in the property market through the mortgage market. In this section, we investigate whether and how the rental market adjusted to the physical event and the central guidance. To study how rents reacted to the shock, we make use of proprietary data from Zoopla, the second most popular property portal in the United Kingdom. We estimate the following dynamic difference-in-differences specification with property fixed effects

Figure 26: Flat vs. Non-Flat Difference in Log Rents of Rental Listings in London

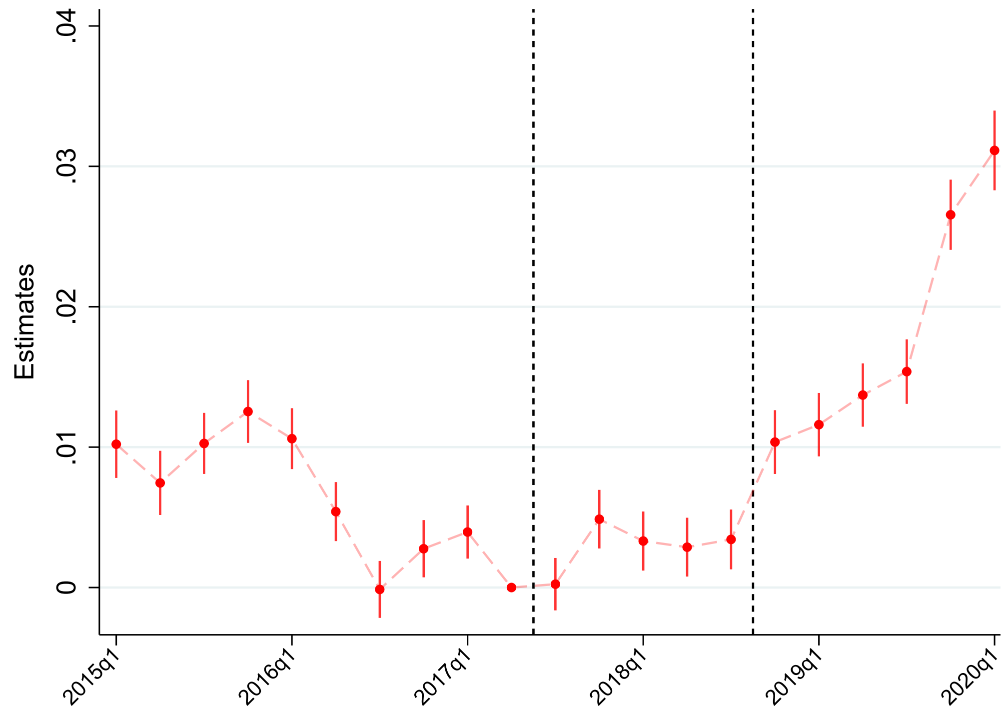

Note: This figure presents difference-in-difference estimates. The plot reports the quarterly difference in mean log rents of rental listings for flats and non-flats along with the 95% confidence interval estimated from a normal distribution. The estimates adjust for property fixed effects. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

for all available rental listings in London:

$$\ln \text{rent}_i = \beta_0 + \sum_{\tau=2015Q1}^{2020Q1} \alpha_\tau \, \mathbb{1}(i \in \text{flat}, t(i) = \tau) + \gamma_{t(i)} + \rho(i) + \varepsilon_i \tag{40}$$

where i indexes each rental listing; $\gamma_{t(i)}$ is an indicator for quarter-year $t(i)$ for listing i; $\rho(i)$ is a fixed effect for property i; and α_τ is the difference-in-differences coefficient for quarter-year τ. Notice that a flat indicator which would be included in a traditional difference-in-differences specification is excluded here since it is absorbed by the property fixed effect. Figure 26 plots estimates of α_τ and its confidence intervals.

There are two patterns worth noting. First, Figure 26 shows that rents did not adjust immediately following the physical event in 2017Q2. In contrast, our preceding analysis shows that prices responded immediately (see Figure 21). This evidence allows us to rule out the possibility that the decrease in prices immediately following the physical event was due to a drop in the willingness to reside in flats arising either due to changes in the perceived quality or risk associated with residing in flats. Second, Figure 26 shows that rents of affected properties adjusted sharply upwards in response to the central guidance by an annualized rate of $\sim$2%. In comparison, the growth of flat prices in London appear to not have dropped sharply following central guidance (see Figure 21).

We interpret the empirical evidence through the lens of a standard user-cost model, which suggests that (i) the post-shock drop in prices and number of mortgages were not generated by a decline in the willingness to reside in flats; and (ii) landlords pass short-run increases in user-costs on to renters through increased rents, but do not adjust rents in response to changes in prices.

3.5 Alternative Interpretation: Demand-Side Channels

While our results related to the sharp drop in mortagage originations following Advice Note 14 are consistent with a contraction in credit supply, our analysis does not rule out the possibility that part of the estimated effects reflects a simultaneous reduction in the willingness to own among mortgage-dependent buyers. Henderson and Ioannides (1983) show that, due to the non-divisible nature of housing, owner-occupiers tend to overinvest in their homes, leaving them underdiversified. Building on this insight, Fu (1991) demonstrates that when housing investment risk or price uncertainty increases, the probability of homeownership declines. Empirically, Turner (2003) and Hilber (2005) document patterns consistent with this mechanism, finding that homeownership rates are lower in markets with greater price volatility.

To the extent that mortgaged buyers are more sensitive to price uncertainty potentially due to being more overinvested in housing, greater perceived housing-market uncertainty could cause them to withdraw more strongly than cash buyers. While such a mechanism could contribute to the observed patterns, it is unlikely to be the primary driver of our results. As I discuss below, the nature and timing of events which followed the shocks are difficult to reconcile with a predominantly demand-based explanation.

First, following the fire, flat prices in London declined steadily for five quarters prior to the announcement of Advice Note 14. This decline represents a significant departure from the pre-fire trend in flat prices, both in levels and relative to the trend for non-flats. Given the salient nature of the fire, it is plausible that buyers of flats would have been aware of the decline in flat prices. Consistent with evidence from Kuchler and Zafar (2019), such a sustained decline could have lowered expected capital gains and increased perceived price uncertainty among buyers. Yet despite this steady decline, aggregate mortgage origination volumes and the share of cash transactions remained broadly stable until the issuance of Advice Note 14, when there was a sharp and discontinuous drop in originations and an increase in the cash share of transactions.[60] If the adjustments in the housing market to the shocks were primarily driven by increasing perceptions of risk or price uncertainty, we would expect a more gradual and sustained response in aggregate mortgage origination volumes immediately following the fire. Instead, the large and sudden response observed after more than a year following the fire is more consistent with a sudden tightening in credit supply linked to how lenders interpreted and responded to the guidance in Advice Note 14.

Second, policy and media discussions surrounding Advice Note 14 tended to focus on how the announcement prompted lenders to condition new lending and refinancing on the availability of EWS1 assessments, effectively cutting mortgage lending on flats without certificates.[61] Third, subsequent official announcements concentrated on restoring lending capacity.[62] This focus on lending capacity suggests that lending potentially played a more important role than demand in driving the observed adjustment.

[60]Origination volumes for home movers and small lenders did adjust downwards immediately following the fire, however, these adjustments were not large enough to have a statistically significant impact on aggregate origination volumes for flats in the five quarters preceding Advice Note 14. Moreover, following the initial drop and prior to Advice Note 14, origination volumes did not continue to decrease for home movers and small lenders despite the sustained reduction in prices during this period. In fact, the point estimates tended to increase during this period, although this increase is not statistically significant.

[61]For instance, in December 2020 the Financial Times published an article titled "England's Cladding Crisis Creates 2 Million 'Mortgage Prisoners'" (https://www.ft.com/content/913cc2ab-7fd5-4d41-a097-df408b4fa57d) .

[62]The government provided guidance in July 2021 that lenders should no longer require EWS1 forms for buildings under 18 metres, aiming to ease sales and remortgaging constraints (see https://www.ft.com/content/7ce70a42-04b2-401c-a838-d91be660e641). Subsequently, in December 2022, the Royal Institution of Chartered Surveyors (RICS) issued updated valuation guidance clarifying when EWS1 forms were necessary in order to "unblock the mortgage crisis" (see https://www.insidehousing.co.uk/news/rics-publishes-new-cladding-valuation-guidance-to-unblock-mortgage-crisis-79346).

Nonetheless, it remains possible that heightened price uncertainty reduced demand for flats among mortgage-dependent buyers. In the absence of mortgage application data, the relative importance of supply and demand channels is difficult to quantify precisely. We acknowledge this shortcoming, while noting that the available evidence suggests that reductions in lending likely played an important role in generating the observed outcomes in the housing market.

3.6 Conclusion

In this paper, we study complementary responses in the mortgage, property and rental markets to an adverse and salient shock to the residential properties in England & Wales. We leverage detailed administrative data on the universe of all residential mortgage and property transactions complemented with property-level data on rents in England & Wales.

Our paper offers a more nuanced perspective to the idea of credit being a "financial accelerator". Whilst we document a decline in property prices and mortgage lending which is most pronounced for information opaque first-time buyers, our results also show that cash-buyers were less sensitive to the shock.

Recognizing that lenders and cash buyers react differently to the same risk signal is crucial. First, it implies that models of housing markets should explicitly include a cash-buyer segment since their exclusion can overstate the sensitivity of housing demand, rents, and prices to shocks. This is especially important given that cash buyers account for approximately one-third of the housing market. Second, it also has distributional implications because when lenders are relatively more elastic than cash buyers, a shock can re-allocate the ownership of affected houses from credit constrained buyers to wealthier cash buyers. Such re-allocation of ownership can have long term effects on the distribution of household wealth since housing is the greatest source of savings and assets for most households.

The re-allocation of ownership need not necessarily be problematic. For instance, it might simply reflect a fall in first-time buyers' willingness to reside in affected properties. Our evidence, however, shows that rents on flats rose after the shock, suggesting that the demand to reside in flats did not fall following the shock. The shock may have thus reshuffled ownership without substantially altering the utilization of the affected dwellings.

Beyond these immediate results, our paper offers an analytical approach for financial stability policymakers to study other, related shocks including those from climate risks. Moreover, it showcases a methodological approach on how these transmission channels can be quantified using available data. By doing so, our paper demonstrates how one can leverage multiple data sources to estimate the effects of shocks to properties. By documenting the appeal but also limitations of the data and methodology we employ, our paper highlights the necessity of collecting data at an appropriate level of granularity to study such effects.

Appendix Chapter 1

Figure A.1: Tenure Shares Over Time in London

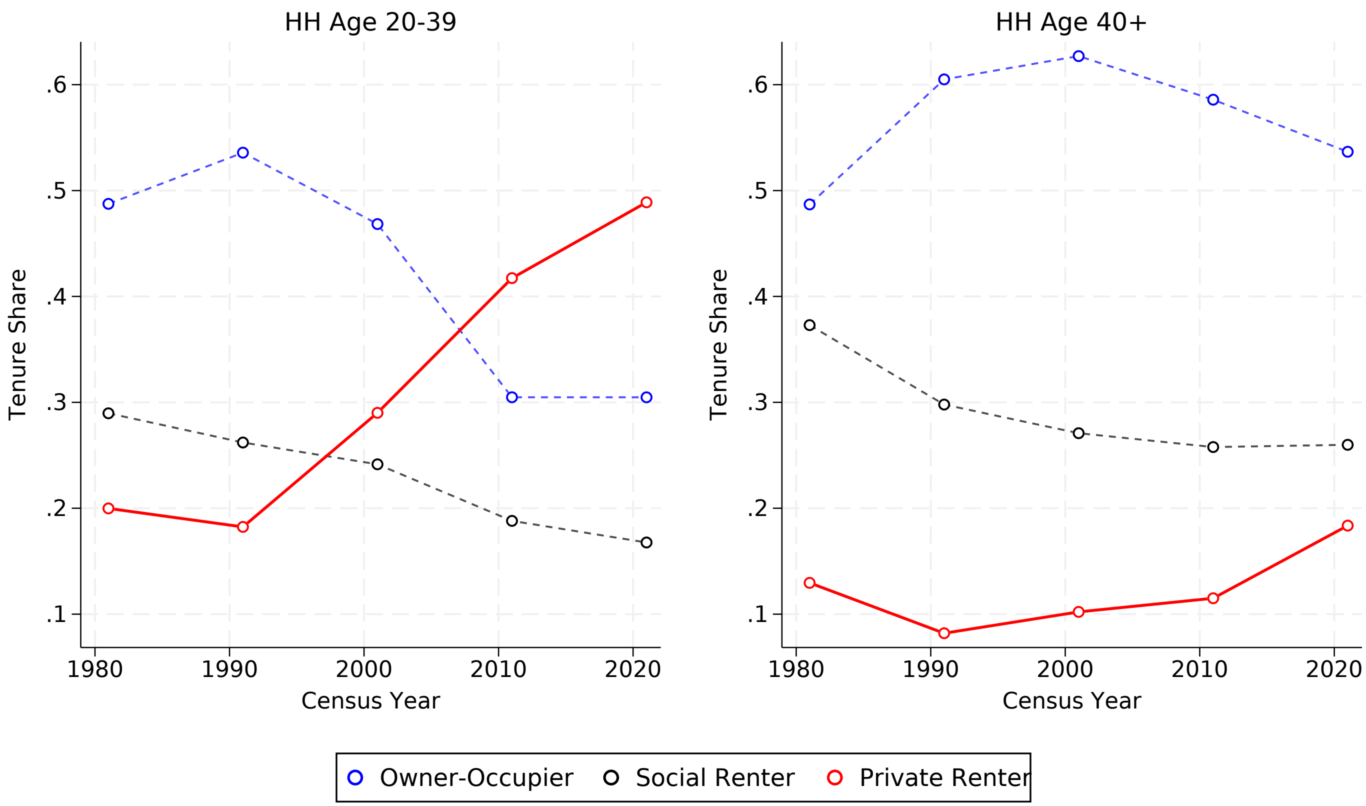

Note: This figure plots the share of households by age in London who belong to one of three housing tenure categories over time. The plots are constructed using Census data.

Figure A.2: Tenure Shares Over Time Outside London

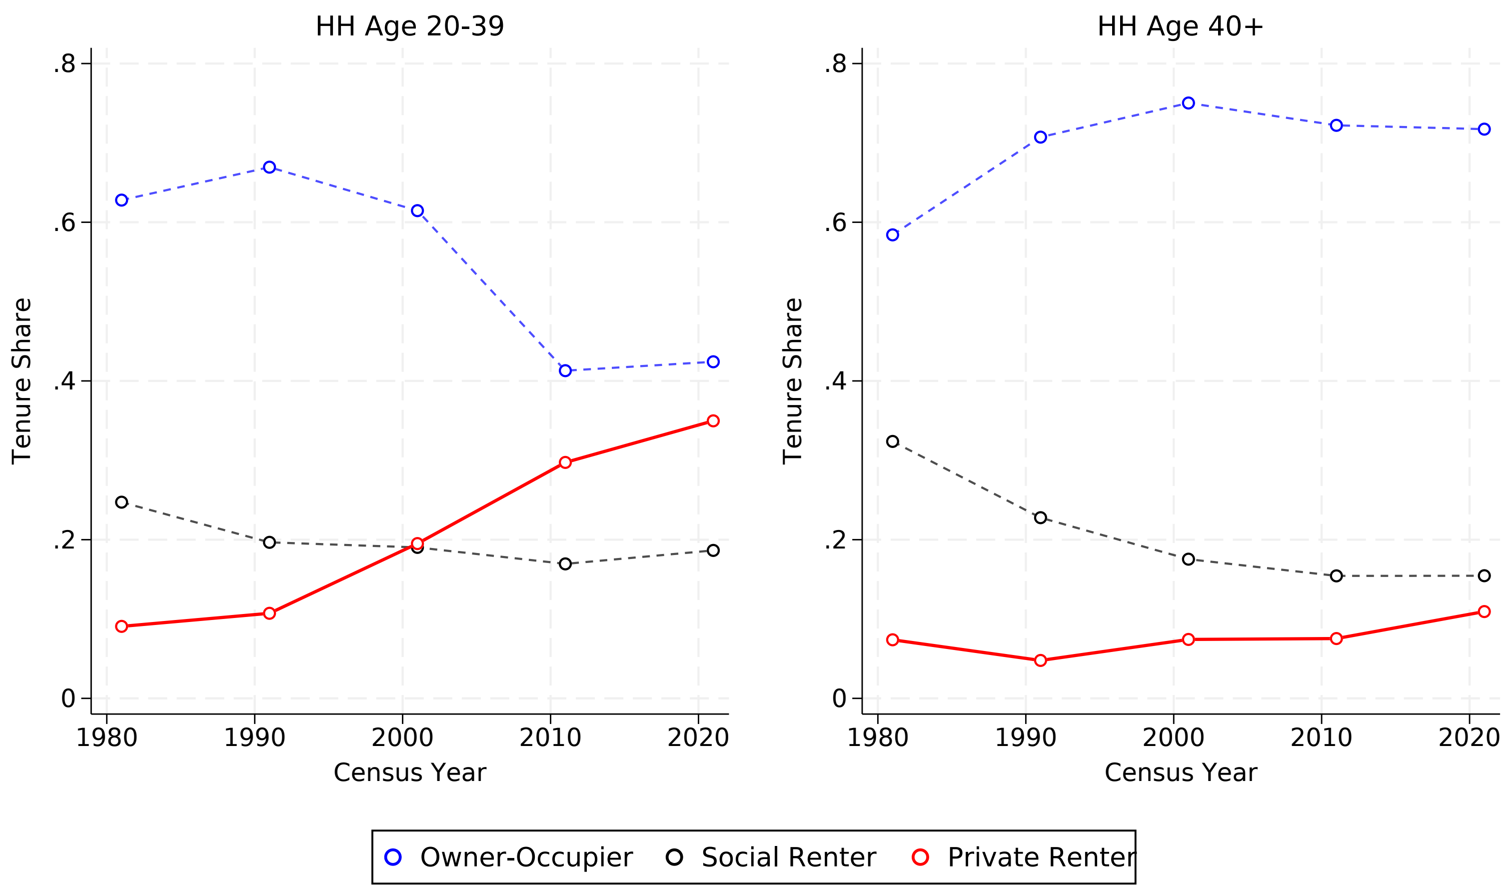

Note: This figure plots the share of households by age outside London who belong to one of three housing tenure categories over time. The plots are constructed using Census data.

A1 Selecting the Polynomial Degree

This section reports the optimal polynomial orders used to estimate three set of regressions in Chapter 1. I use five-fold cross-validation to select the optimal degree in each case. I restrict the search to polynomial degrees one through five. To improve numerical stability, I use orthogonal polynomials.

Table A.1: Optimal Polynomial Order Selected Using Five-fold Cross-Validation For the Mapping of 2014 Rent Rank to 2014 Sale Rank

City	Order
London	4
Cambridge	3
Manchester	4
Birmingham	5
Liverpool	3
Bournemouth	5
Bristol	2
Leeds	5
Newcastle	3
Nottingham	5
Sheffield	4

Mapping Rent Ranks to Sale Ranks—Table A.1 reports the optimal polynomial order by city for the quantile regressions used to estimate the mapping $g_c(q)$ between the 2014 rent rank and 2014 price rank.

Mapping 2014 Ranks to Ranks in Subsequent Years—Table A.2 reports the optimal polynomial degree by city for the quantile regressions used to estimate the mappings m_{rcy} and m_{pcy} between ranks in 2014 and subsequent years. The column on the left

Table A.2: Optimal Polynomial Order Selected Using Five-fold Cross-Validation For the Mapping of 2014 Rank to Future Rank

City	Order	
	Rent	Sale
London	4	3
Cambridge	5	1
Manchester	5	1
Birmingham	5	5
Liverpool	5	2
Bournemouth	5	3
Bristol	2	1
Leeds	3	4
Newcastle	5	3
Nottingham	4	2
Sheffield	5	3

labelled "Rent" presents the optimal polynomial degree for the mapping of ranks for rental listings. The column on the right labelled "Sale" presents the optimal polynomial degree for the mapping of ranks of sale listings.

Notice, although the mappings m_{rcy} and m_{pcy} in Section 1.2.1 are estimated separately for each city-year combination, the reported polynomial degree in Table A.2 varies only by city. This is because to determine the appropriate polynomial degree, I conduct the cross-validation once for each city-tenure combination by pooling data from all years for each city-tenure pair. Specifically, for each city and tenure (rental or sale), I estimate pooled regressions $q^r_{icy} = m_{rc}(q^r_{ic2014})$ and $q^p_{icy} = m_{pc}(q^p_{ic2014})$ that pool data from all subsequent years $y > 2014$, and select the polynomial degree that minimizes the cross-validation error for the regressions for each city-tenure pair. Table A.2 reports the polynomial orders from this procedure.

When estimating the mappings m_{rcy} in Section 1.2.1, I use the polynomial order in the rent column for city c in Table A.2. Similarly, the mappings m_{pcy} use the polynomial order in the sale column for city c. This procedure assumes that while the mapping between ranks can shift over time, its functional form i.e., the degree of the polynomial, remains stable within a city for each tenure.

Table A.3: Optimal Polynomial Order Selected Using Five-fold Cross-Validation For the Relationship Between Rent-to-Price and Price

City	Order
London	2
Cambridge	5
Manchester	5
Birmingham	2
Liverpool	5

Estimating the Rent-to-Price Ratio as a Function of Price—Table A.3 reports the optimal polynomial degree by city for the OLS regressions used to estimate the relationship between the rent-to-price ratio and price. Section 1.3.2 discusses these regressions and how they are used as inputs to construct the lowess plots in Figure 8. Table A.3 reports the polynomial orders for the five cities that appear in the lowess plots in Figure 8.

Figure A.3: Evolution of Property Type Rankings in the Rent and Price Distributions

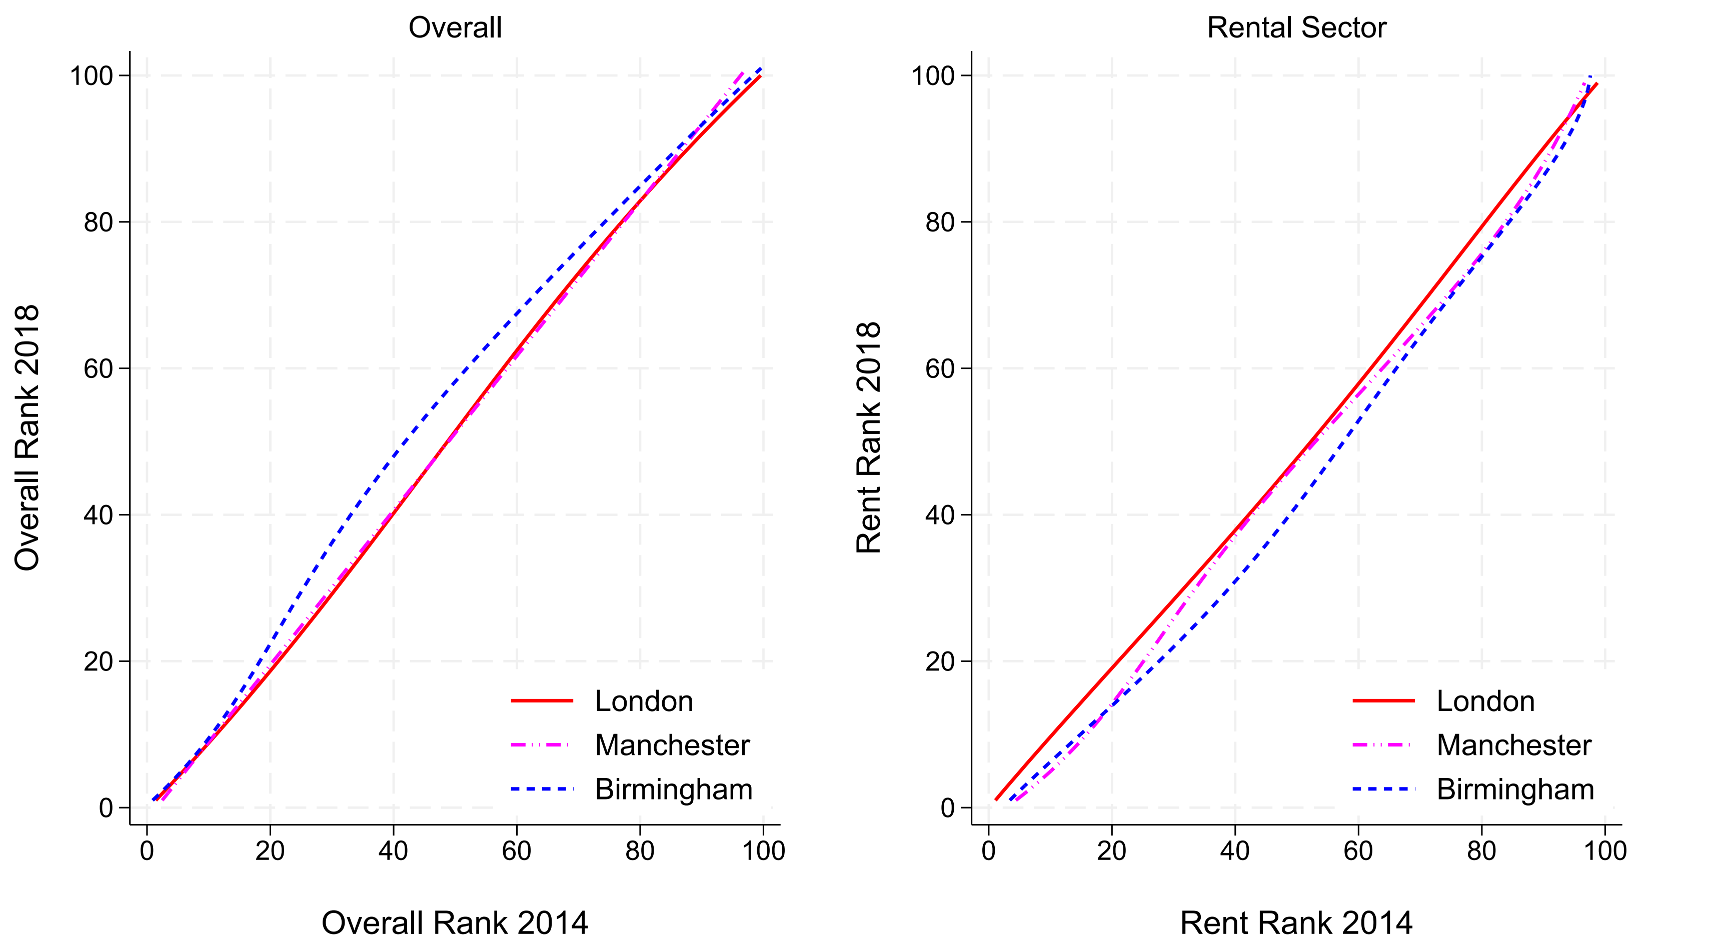

Note: This figure plots the average mapping between the quality ranks of a property in 2014 and 2018 by city. The figure on the left presents the mapping for the entire housing market whereas the figure on the right presents the mapping for the rental sector.

Figure A.4: Rent-to-Price Ratio and Capital Gains (2014-2018)
Rent-to-Price Imputed using Polynomial Regression

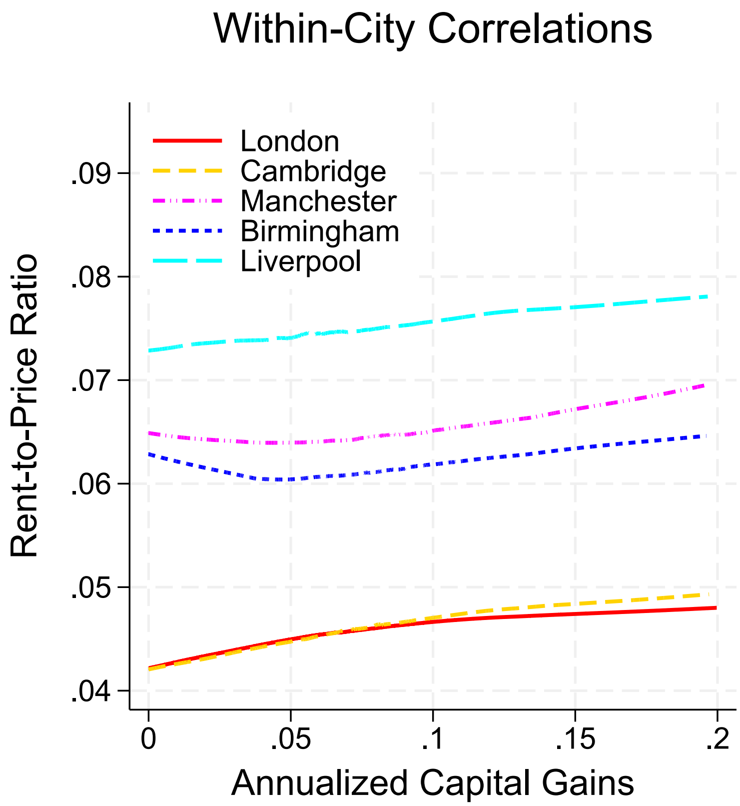

Note: This figure plots the relationship between the rent-to-price ratio and capital gains within cities. The capital gains are computed using data on properties that were sold both in 2014 and 2018. The rent-to-price ratio is imputed for each property in the repeat sales sample given the 2014 sales price. The imputation is based on the same polynomial regressions used to impute rent-to-price ratios in the right plots of Figure 8. The figure presents lowess-smoothed plots for each city of the conditional expected rent-to-price ratio for properties with a given level of annualized capital gains.

Figure A.5: Rent-to-Price Ratio and Capital Gains (2010-2018)
Rent-to-Price Imputed Using Regression Specification That is Linear in Prices

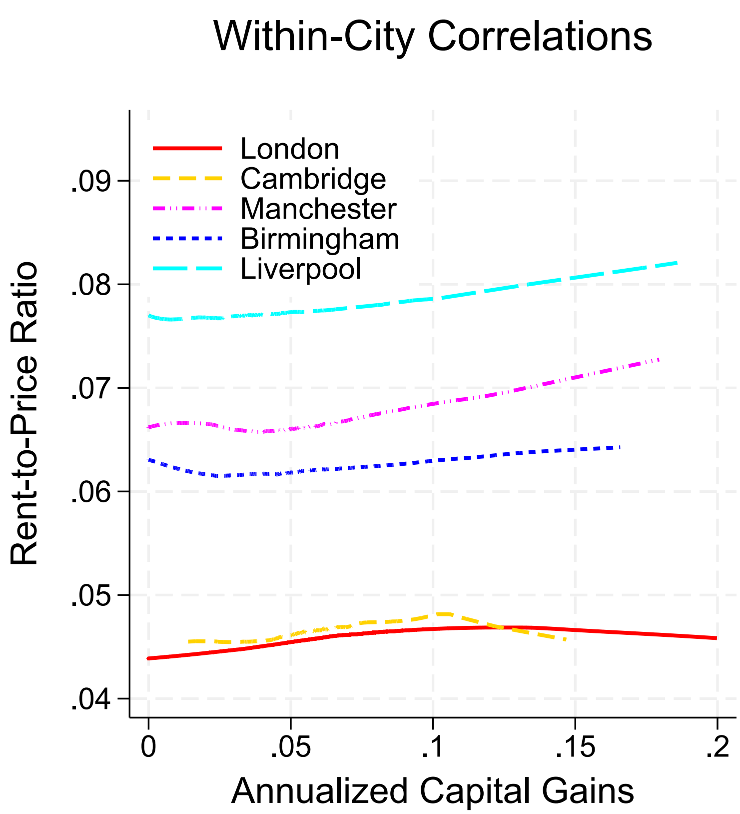

Note: This figure plots the relationship between the rent-to-price ratio and capital gains within cities. The capital gains are computed using data on properties that were sold both in 2010 and 2018. The rent-to-price ratio is imputed for each property in the repeat sales sample given the 2010 sales price. The imputation of rent-to-price is conducted identically to the right plots of Figure 8 with the exception that the regression specifications used are linear in prices. The figure presents lowess-smoothed plots for each city of the conditional expected rent-to-price ratio for properties with a given level of annualized capital gains.

Appendix Chapter 2

B.1 Landlord Marginal Profit From Operating in Higher Quality Segments Increasing in a

This section verifies that the cross-derivative of the landlord's profit function is positive as in Equation 20. Recall that the variable cost $\mu_{ct}(h,a)$ depends on a through the LTV function $\phi^a_{ct}(p_{cth})$. Equation 20 depends on a first derivative and a cross-derivative of $\mu_{ct}(h,a)$. Using the chain rule, we can write these derivatives as:

$$\frac{\partial \mu_{ct}(h,a)}{\partial a} = \frac{\partial \mu_{ct}(h,a)}{\partial \phi}\frac{\partial \phi^a_{ct}(p_{cth})}{\partial a}$$

$$\frac{\partial^2 \mu_{ct}(h,a)}{\partial h \partial a} = \frac{\partial^2 \mu_{ct}(h,a)}{\partial \phi^2}\frac{\partial \phi^a_{ct}(p_{cth})}{\partial h}\frac{\partial \phi^a_{ct}(p_{cth})}{\partial a} + \frac{\partial \mu_{ct}(h,a)}{\partial \phi}\frac{\partial^2 \phi^a_{ct}(p_{cth})}{\partial h \partial a}$$

Using the fact that $\phi^a_{ct}(p_{cth}) = 1 - a/p_{ct}(h)$, we can rewrite the first and second derivatives as follows:[63]

$$\frac{\partial \mu_{ct}(h,a)}{\partial a} = -\frac{\partial \mu_{ct}(h,a)}{\partial \phi}\frac{1}{p_{ct}(h)} \tag{B.1}$$

$$\begin{aligned}\frac{\partial^2 \mu_{ct}(h,a)}{\partial h \partial a} &= -\frac{\partial^2 \mu_{ct}(h,a)}{\partial \phi^2}\frac{a \times \partial p_{ct}(h)/\partial h}{p_{ct}(h)^3} + \frac{\partial \mu_{ct}(h,a)}{\partial \phi}\frac{\partial p_{ct}(h)/\partial h}{p_{ct}(h)^2} \\ &= \frac{\partial p_{ct}(h)/\partial h}{p_{ct}(h)^2}\left(-\frac{\partial^2 \mu_{ct}(h,a)}{\partial \phi^2}\frac{a}{p_{ct}(h)} + \frac{\partial \mu_{ct}(h,a)}{\partial \phi}\right)\end{aligned} \tag{B.2}$$

[63] Recall that $\phi^a_{ct}(p_{cth})$ does not depend directly on h. h only influences the loan-to-value ratio because prices p_{cth} vary with h.

Equation B.2 shows that the cross-derivative of variable costs $\frac{\partial^2 \mu_{ct}(h,a)}{\partial h \partial a}$ is positive whenever $\frac{\partial \mu_{ct}(h,a)}{\partial \phi} - \frac{\partial^2 \mu_{ct}(h,a)}{\partial \phi^2} \frac{a}{p_{ct}(h)} > 0$, and negative otherwise. The sign of this difference depends on the values of a and $p_{ct}(h)$. For instance, when the variable cost is quadratic, the cross-derivative of variable costs is positive whenever $a/p_{ct}(h) < 1/2$, and negative otherwise.[64]

Finally, substituting the derivatives in Equations B.1 and B.2 into the cross-derivative in Equation 20 yields:

$$\begin{aligned}\frac{\partial^2 \Pi_{ct}(h \mid a)}{\partial h \partial a} &= \frac{\partial p_{ct}(h)/\partial h}{p_{ct}(h)} \left(\frac{\partial \mu_{ct}(h,a)}{\partial \phi} + \frac{\partial^2 \mu_{ct}(h,a)}{\partial \phi^2} \frac{a}{p_{ct}(h)} - \frac{\partial \mu_{ct}(h,a)}{\partial \phi} \right) \\ &= \frac{\partial p_{ct}(h)}{\partial h} \times \frac{\partial^2 \mu_{ct}(h,a)}{\partial \phi^2} \times \frac{a}{p_{ct}(h)^2} > 0 \end{aligned} \tag{B.3}$$

Notice that the first derivative $\frac{\partial \mu_{ct}(h,a)}{\partial \phi}$ drops out of this expression. The remaining three terms in the product in Equation B.3 are all positive. The first term is positive because prices increase monotonically in h. The second term is positive because variable costs are convex in the LTV ϕ. Finally, the last term is positive because both a and p are positive. This verifies that the cross-derivative of the profit function is positive.

I showed above that marginal profits from operating in higher quality segments is increasing in a even though the marginal costs from operating in higher quality segments can both increase or decrease in a. This difference arises because the cross derivative of cost depends on the difference $\frac{\partial \mu_{ct}(h,a)}{\partial \phi} - \frac{\partial^2 \mu_{ct}(h,a)}{\partial \phi^2} \frac{a}{p_{ct}(h)}$, which may

[64] To see this, consider a simple case where the variable cost takes the quadratic form $\mu_{ct}(h,a) = \gamma \phi^a_{ct}(p_{cth})^2$, the first derivative is $2\gamma\phi^a_{ct}(p_{cth})$ and the second derivative is 2γ. In this case, the cross-derivative $\frac{\partial^2 \mu_{ct}(h,a)}{\partial h \partial a}$ is positive whenever $2\gamma\phi^a_{ct}(p_{cth}) - 2\gamma\frac{a}{p_{ct}(h)} > 0$ which occurs when $\frac{a}{p_{ct}(h)} < 1/2$.

be either positive or negative. In contrast, the cross-derivative of the profit function does not inherit this switching behaviour, because it depends only on the level of the second derivative $\frac{\partial^2 \mu_{ct}(h,a)}{\partial \phi^2}$, not on its difference relative to the first derivative. Since the variable cost is convex in the loan-to-value ratio, the second derivative is positive. It follows that the cross-derivative of profits is also positive. The landlord's single-crossing condition is therefore always satisfied, even though the marginal cost is not monotonic in a.

B.2 Quality Segmentation and the Usercost-to-Rent Ratio

The pattern of quality segmentation observed in the data implies a specific sorting of households into tenure and quality across cities. Why do households sort in this manner? In this section, I show that this pattern of sorting is explained by differences in the relative costs of owning and renting across cities.

Figure B.1 plots the monthly unobserved user cost of owner-occupation in levels and in comparison to rents. While user costs increase convexly with housing quality across all cities, the increase is dramatically more pronounced in London, reflecting disparities in the price distribution. Since differences in rents are more muted, this generates a corresponding divergence in the user costs-to-rent ratio between London and the comparison cities.

The demand model implies that these differences in relative costs should generate a wedge between households' optimal choice of quality by tenure. In partic-

Figure B.1: Estimated Usercosts of Owner-Occupation in Levels and Relative to Rents 2015

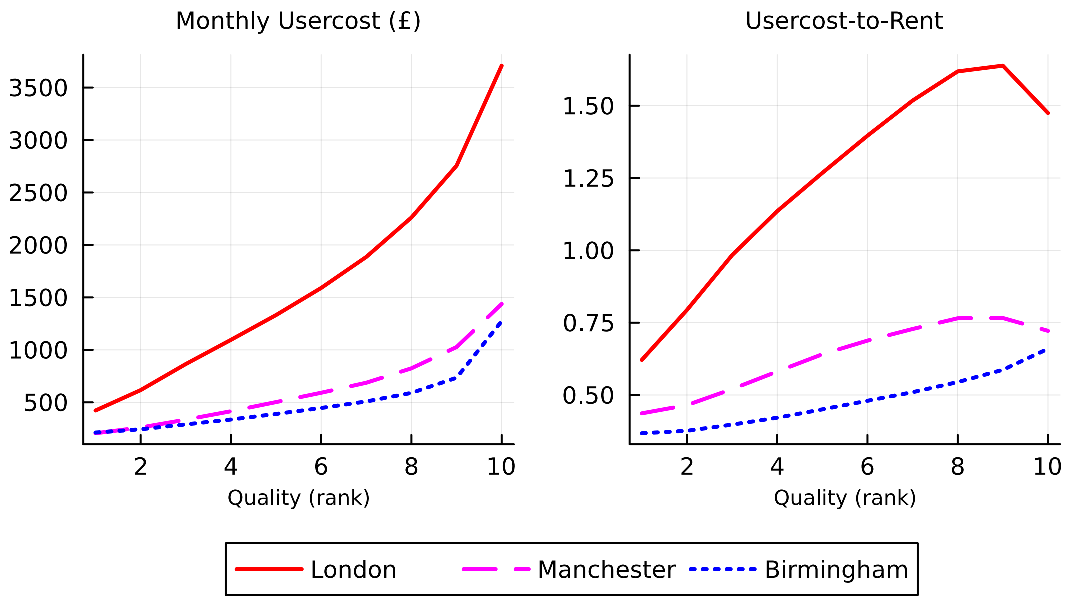

Note: This figure plots estimated unobserved usercosts of owner-occupation and the ratio of usercosts-to-observed rents by quality across cities in 2015

ular, households in London should prefer more quality when they rent compared to when they own given the steep growth in the relative cost of owning. Figure B.2 shows that this is indeed the case. The figure plots a weighted average of the optimal quality preferred by households when they do not face credit constraints by tenure at different levels of income. The weighted average is taken over unobserved types and household assets. The figure shows that households in London prefer more quality when renting across almost all levels of income, and that the difference becomes more pronounced as income increases. The pattern is reversed in Birmingham where households prefer more quality when owning across the entire income range.

Manchester is an intermediate case where the preference for quality when renting is greater only at higher incomes.

Figure B.2: Average Unconstrained Optimal Quality by Tenure and Income

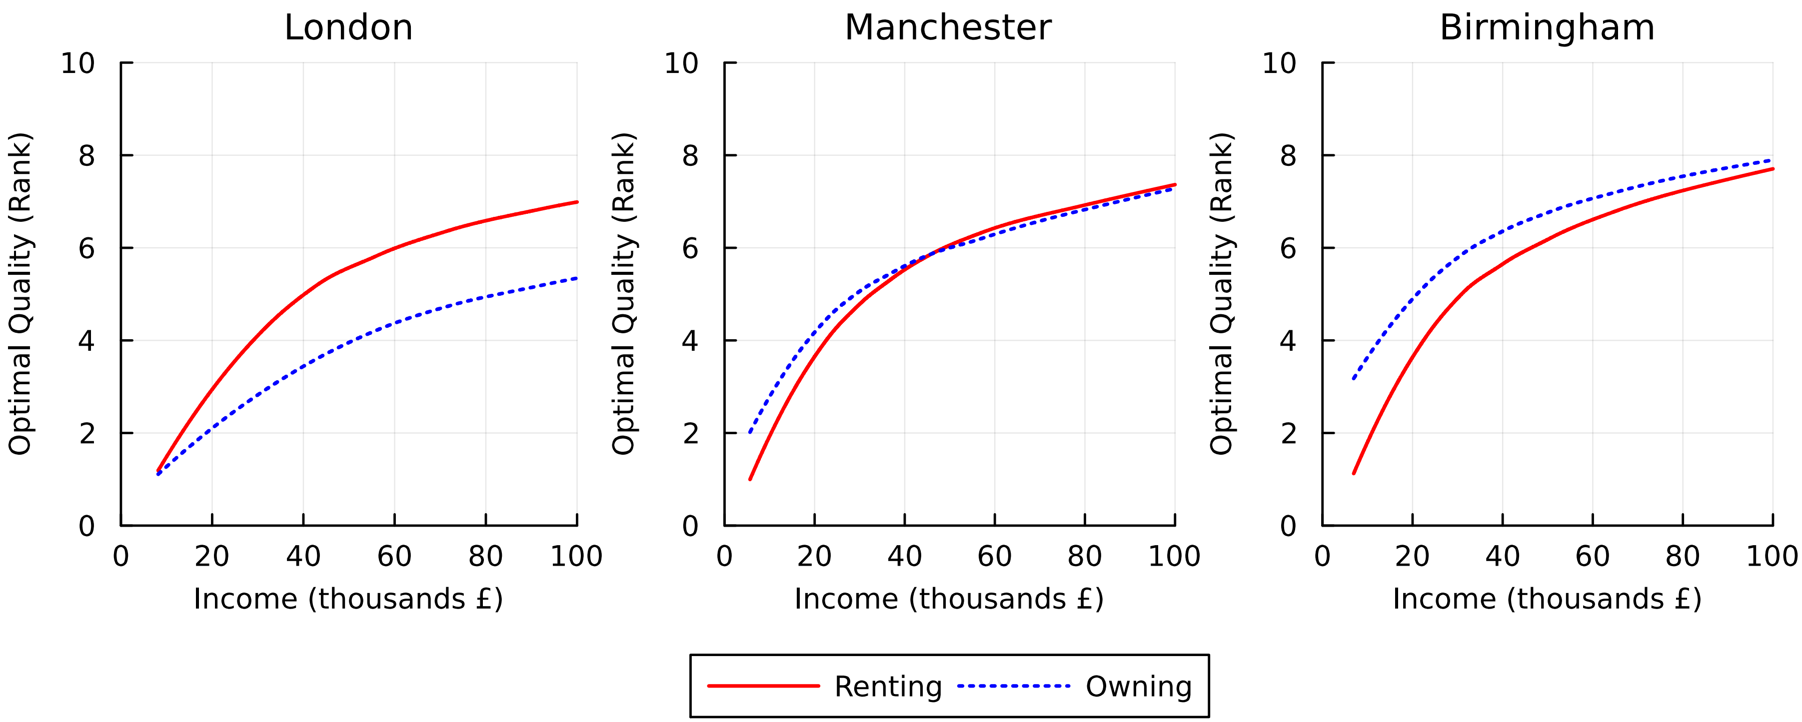

Note: This figure plots estimated income-specific means of the optimal quality a household would choose to rent or own in the absence of credit constraints. The estimates are based on a local regression of optimal quality against household income.

B.3 Model Fit

Figure B.3: Ownership Rates (2015)

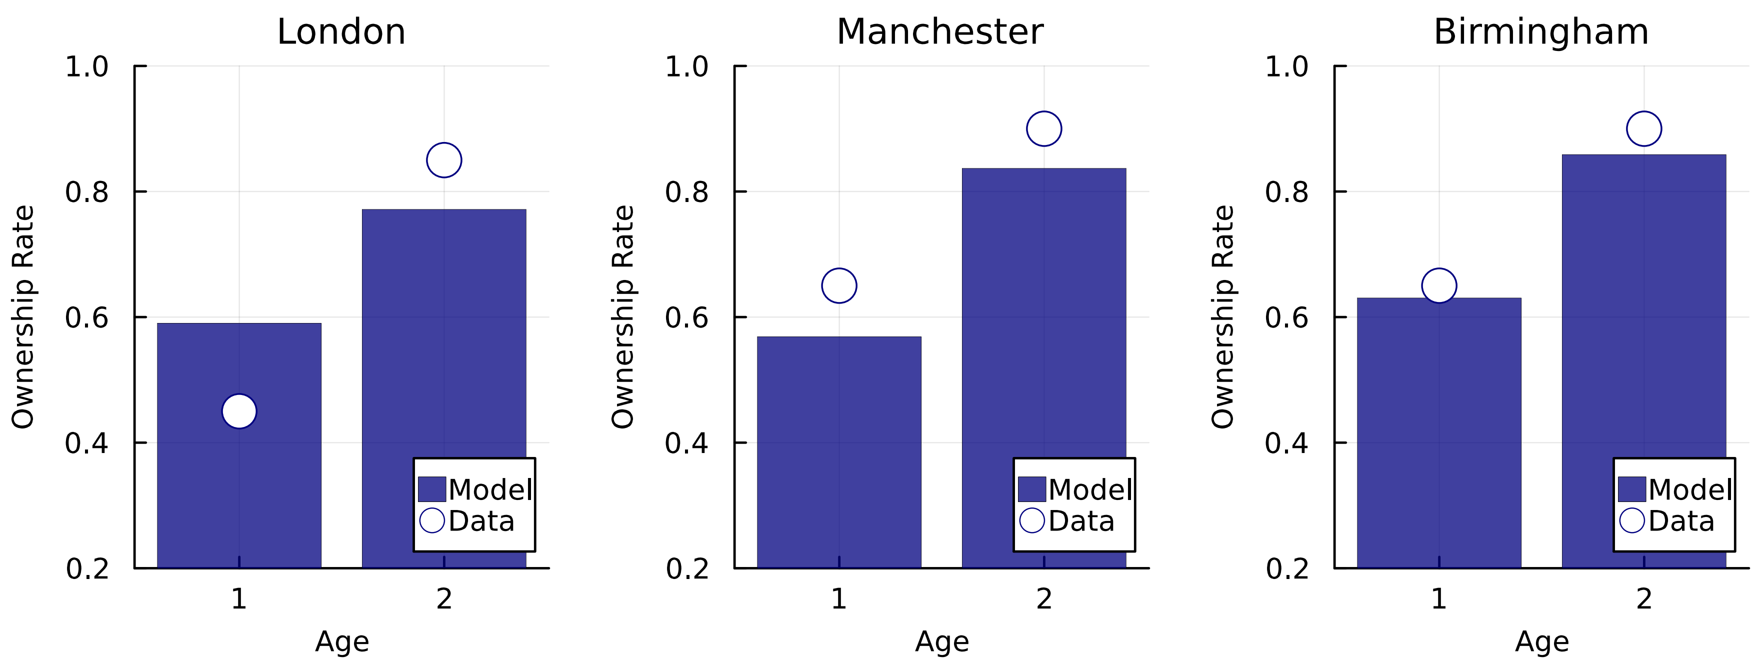

Note: This figure compares the ownership rates implied by the estimated model against ownership rates in the data by city and age group in 2015.

Figure B.4: Aggregate Cumulative Demand and Supply (2015)

(a) Overall Housing Market

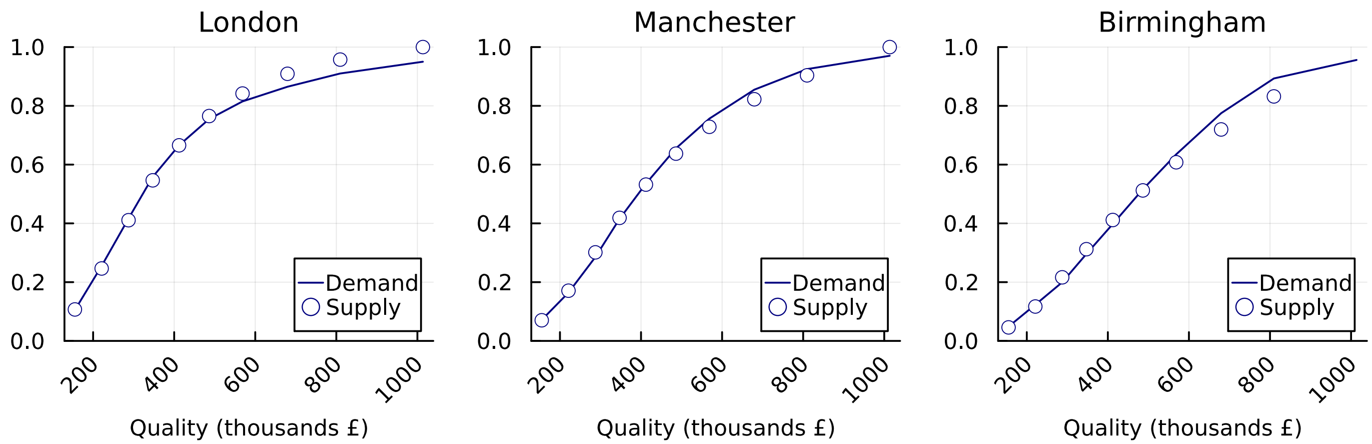

(b) Market for Rental Housing

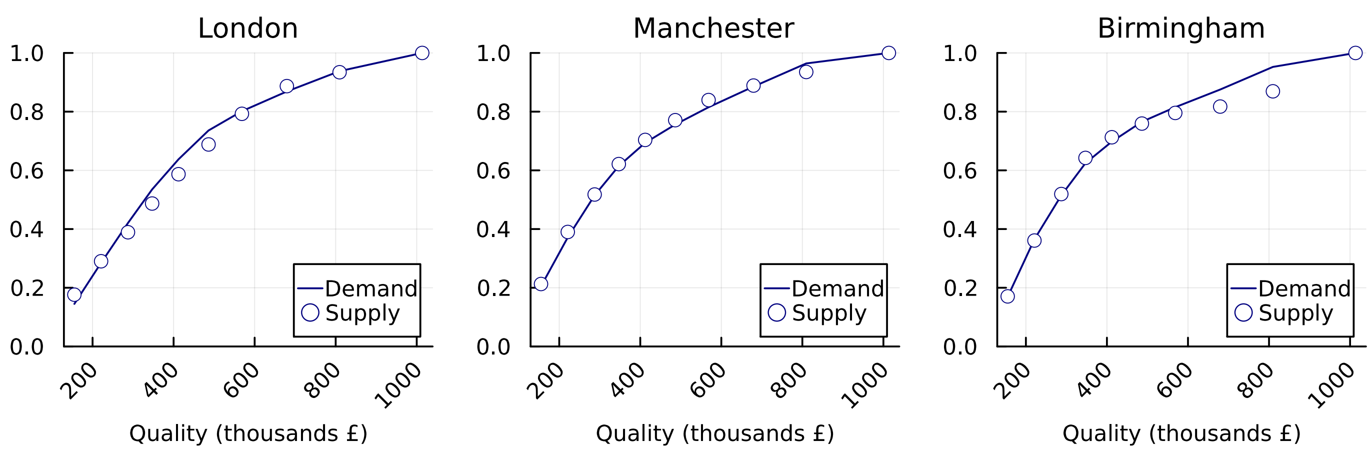

Note: This figure presents the cumulative demand and supply functions by quality at equilibrium based on the estimated model. The vertical axis measures the cumulative share for each quality.

Figure B.5: Distribution of Equilibrium Prices and Rents (2015)

(a) House Price

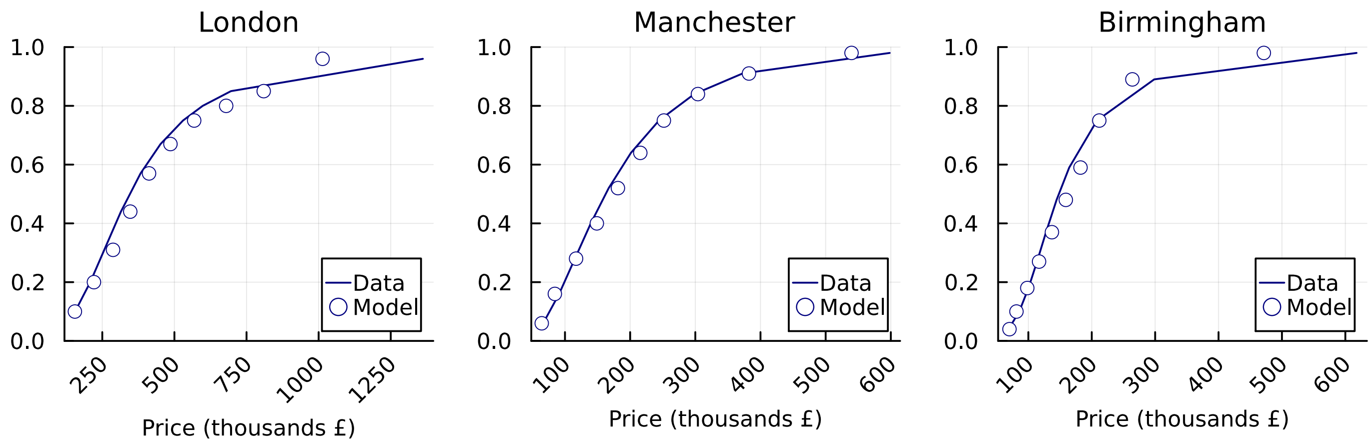

(b) Monthly Rents

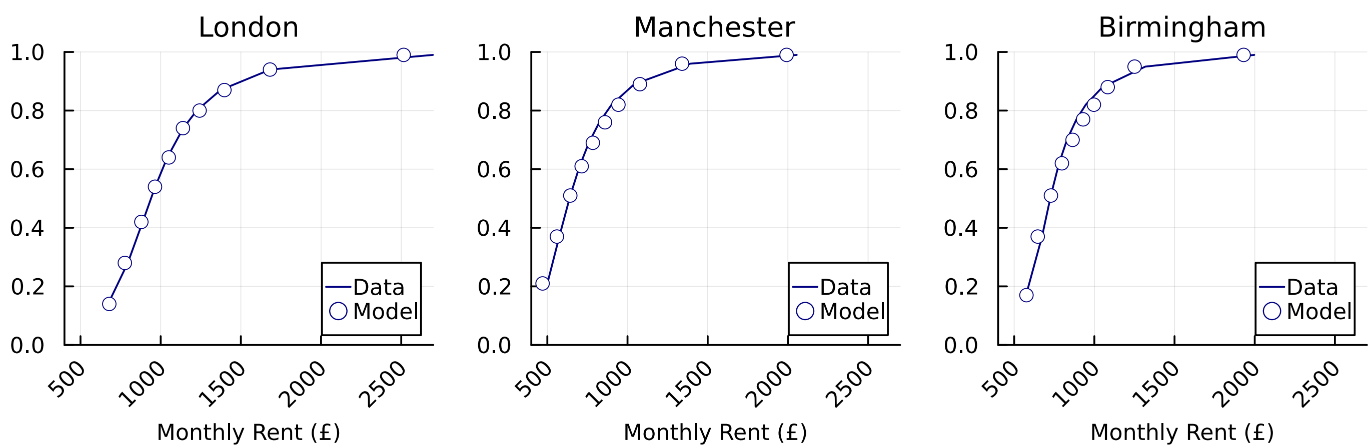

Note: This figure compares the cumulative distribution of equilibrium prices and rents implied by the estimated model to the distribution in the data. The solid curves labelled Data plot the empirical CDFs obtained by fitting the GB2 distribution. Each scatter point corresponds to one of ten quality levels in the model. The height of the price and rent scatters depends on the quality distribution of properties in the overall and rental markets respectively.

Figure B.6: Landlord Assignment Asset Cutoffs by House Quality (2015)

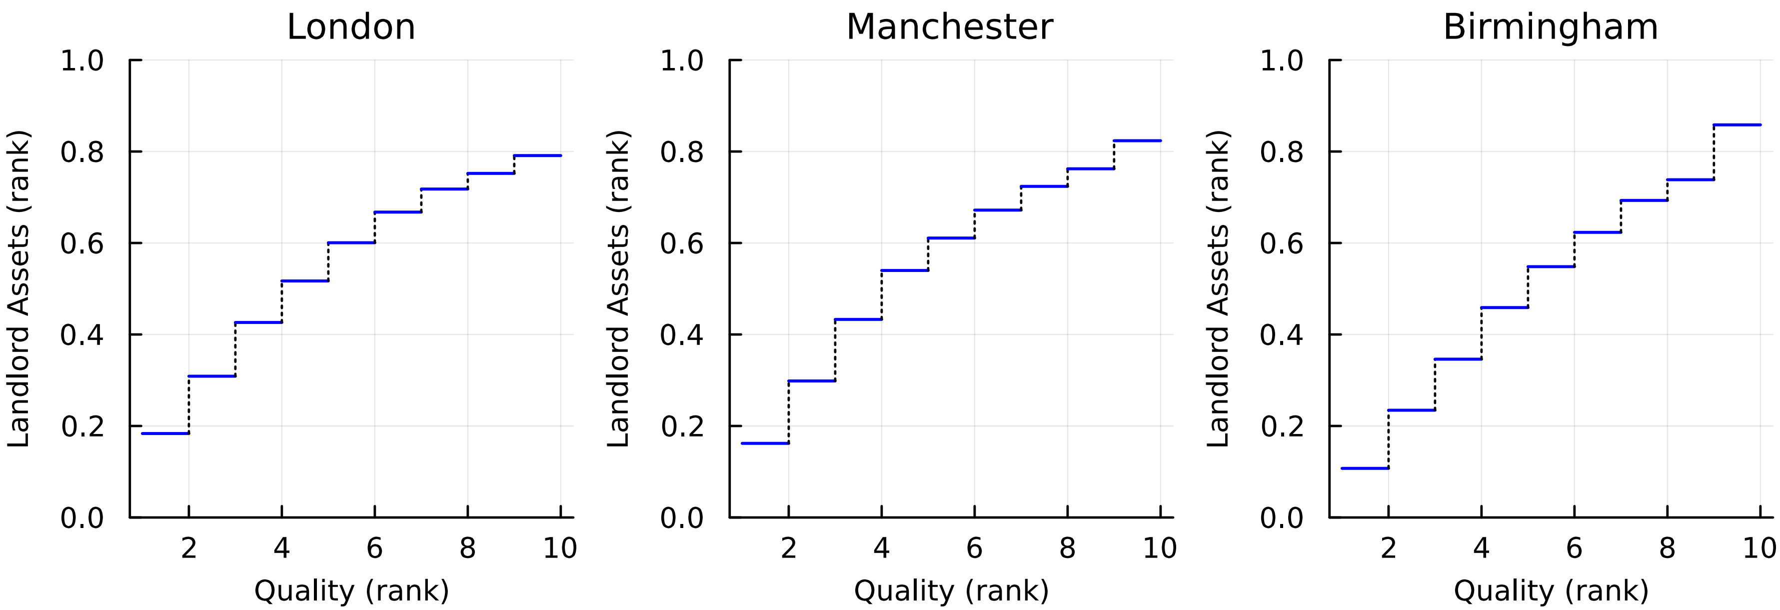

Note: This figure presents the nine asset cutoffs which define the assignment of landlords to house quality. The leftmost step represents the asset threshold which makes landlords indifferent between operating in the first and second quality levels. The rightmost step represents the asset indifference threshold between the ninth and tenth quality levels.

Appendix Chapter 3

Table C.1: Summary statistics of key variables in our data sets

	Mean	Median	St.Dev.
Mortgage transactions (PSD)			
Price £	295,002	235,000	373,310
First time buyer	0.477	0	0.499
Home mover	0.523	1	0.499
Number of rooms	3.07	3	3.88
Number of transactions	3,114,559		
Property transactions (PPD)			
Price £	289,947	225,000	317,890
Number of transactions	4,639,365		
Rental listings (Zoopla)			
Rent	467	375	340.15
Number of observations	4,470,104		

Note: This table presents summary statistics for the three datasets used in the analysis. The statistics for PPD and PSD are based on data for England & Wales between 2015 and 2020. The statistics for Zoopla is based on data for listings in London.

Figure C.1: Flat vs. non-Flat Difference in Log Number of Originations in England & Wales

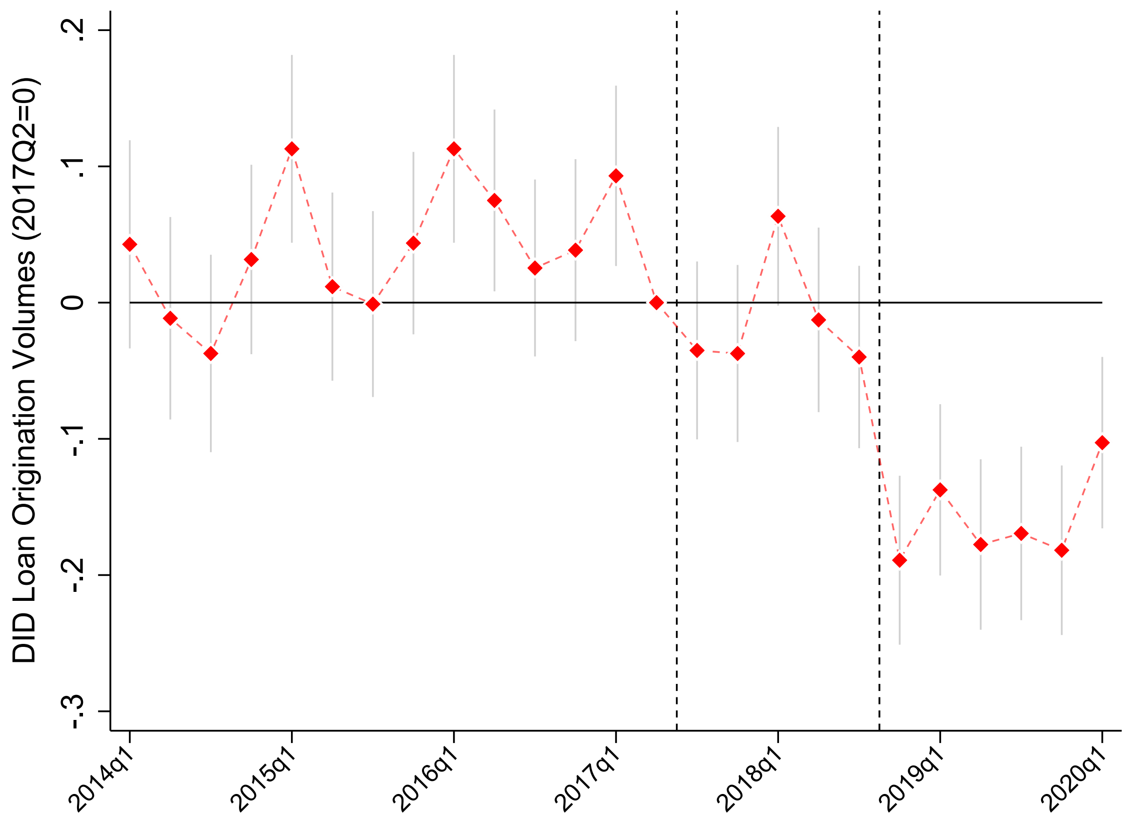

Note: This figure presents difference-in-difference estimates from a Poisson regression. The plot reports the quarterly difference in mean log number of flats and non-flats along with the 95% confidence intervals. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

C.1 Less credit to highly leveraged borrowers?

Our main results suggested that originations of mortgages against affected properties fell sharply and in a sustained manner after the central guidance, with the largest drops occurring in London. This might have further reduced demand for affected properties if it was mainly already credit-constrained borrowers affected by a tightening in underwriting standards. In this section, we explore whether the impact of the shock to origination volumes varied by borrower leverage.

Specifically, we examine responses in origination volumes by loan-to-income (LTI) and loan-to-value (LTV) ratios. The first ratio, LTI, is a measure of initial affordability of a mortgage. The shock, in particular the central guidance, increased insurance costs for affected properties as well as costs of maintenance and repairs (FCA, 2022; DLUHC, 2022). It might have changed lenders' perceptions about the risk of non-repayment due to an increase in non-mortgage costs associated with the ownership of flats. The second ratio, LTV, is a measure of borrower leverage. The shock reduced prices of affected properties and increased uncertainty around the true value. This might have changed lenders' willingness to lend against high-LTV mortgages following the physical event and central guidance.

Loan-to-Income Ratio (LTI): To explore heterogeneity in origination by LTI, we estimate Equation 38 for three different LTI bands. For the first regression, we restrict the sample to only include mortgage originations with LTI between 0-4, i.e. those mortgages that are least risky. We define the count Q_{iflt} in Equation 38 using only mortgage originations in this sample—this count measures the number

Figure C.2: Flat vs. Non-Flat Difference in Log Number of Mortgage Originations in London, by Borrower LTI

(a) Low LTI: 0 to 4 (b) Medium LTI: 4 to 4.5 (c) High LTI: $\geq$ 4.5

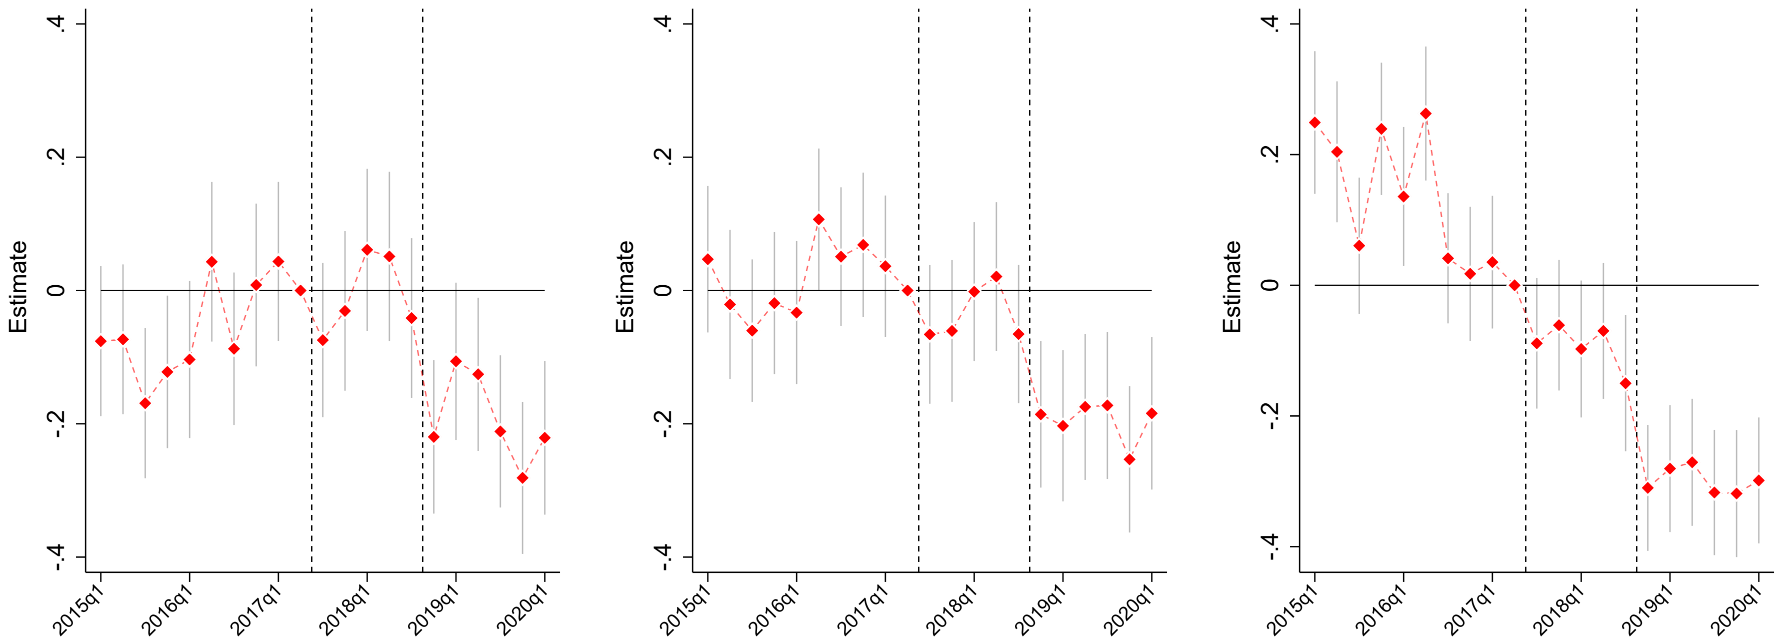

Note: This figure presents difference-in-difference estimates from three different Poisson regressions. Each figure plots the quarterly difference in mean log number of mortgages originated for flats and non-flats in England & Wales by borrowers with different LTIs along with the 95% confidence intervals. The estimates adjust for postcode district fixed effects. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

of originations with LTI between 0-4. We similarly conduct this exercise for a medium LTI band ($4 \leq \text{LTI} \leq 4.5$) and a high LTI band (LTI> 4.5), i.e. those mortgages that are riskiest. For each band, we define the count variable to only include originations of mortgages within the relevant LTI band.

Figure C.2 plots the difference-in-differences coefficients from these regressions. We observe that the contraction in lending in response to the central guidance is increasing in the LTI. The point estimates indicate that originations for mortgages on flats with LTIs below 4.5 fell by around $\sim$20%. By contrast, mortgages on flats with higher LTIs experienced a larger contraction in originations of $\sim$30%. This difference of $\sim$10 percentage points in the point estimates is consistent with the central guidance increasing lenders' concerns that flat owners would be unable to service their monthly mortgage payments. However we should be cautious in interpreting this difference as the estimates are noisy and do not appear to be statistically significant as judged by the overlap of the confidence intervals.

LOAN-TO-VALUE RATIO: We next study heterogeneity in the evolution of the number of mortgages against flats across different LTV bands. Analogously to our earlier analysis of heterogeneity across LTI bands, we re-estimate Equation 38 for three different LTV bands.

For the first regression, we restrict the sample to only include mortgage originations with LTV between 0% and 75%, i.e. those mortgages that are least risky. We define the count Q_{iflt} in Equation 38 using only mortgage originations in this sample—this count measures the number of originations with LTV between

Figure C.3: Flat vs. Non-Flat Difference in Log Number of Mortgage Originations in London, by Borrower LTV

(a) Low LTV: 0% to 75%　　(b) Medium LTV: 75% to 90%　　(c) High LTV: 90% to 100%

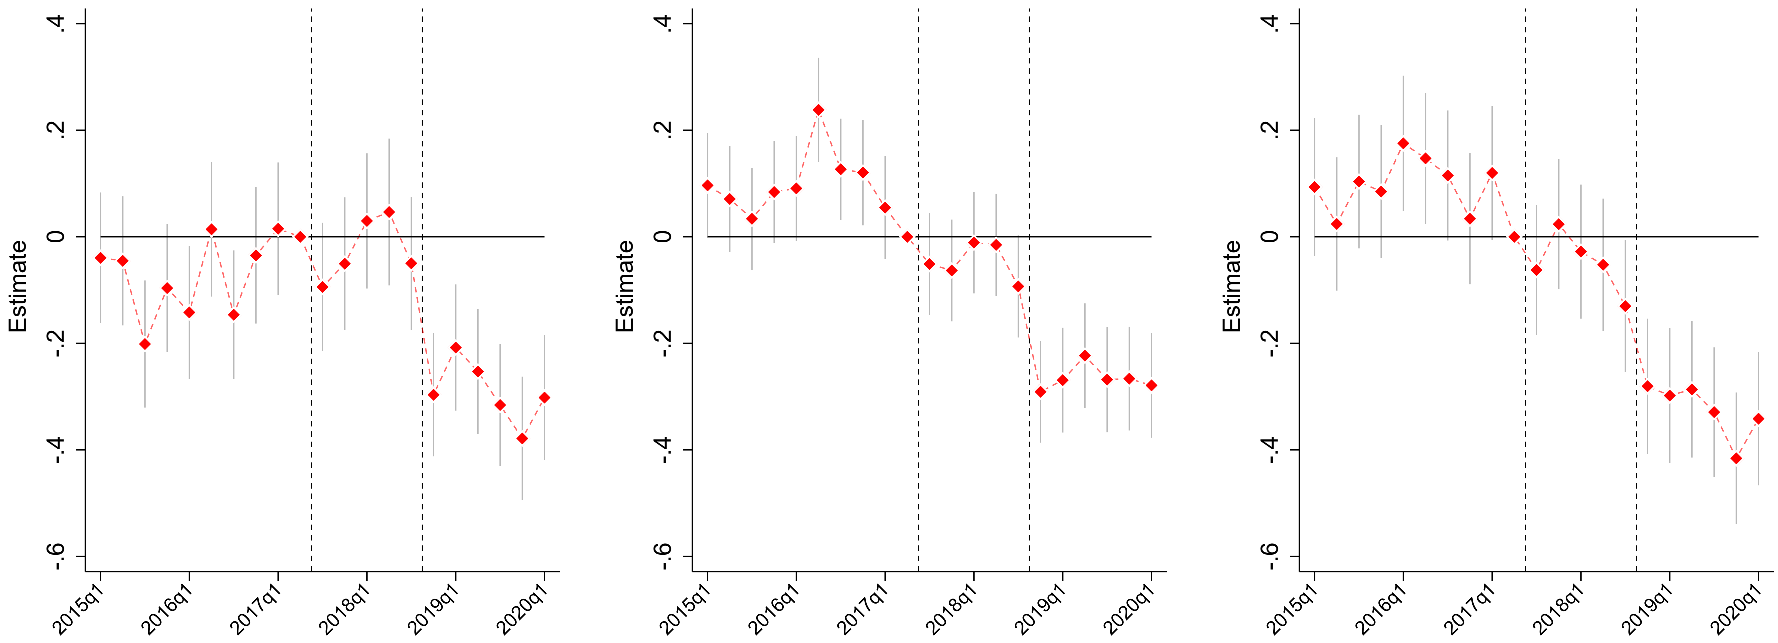

Note: This figure presents difference-in-difference estimates from three different Poisson regressions. Each figure plots the quarterly difference in mean log number of mortgages originated for flats and non-flats in London by borrowers with different LTVs along with the 95% confidence intervals. The estimates adjust for postcode district fixed effects. All estimates are relative to 2017Q2. The figure uses vertical dashed lines to mark two events: the leftmost vertical dashed line represents the period of the physical event whereas the rightmost vertical dashed line represents the central guidance.

0% and 75%. We similarly conduct this exercise for a medium LTV band (75% $\leq$ LTV $\leq$ 90%) and a high LTV band (LTV$>$ 90%), i.e. those mortgages that are riskiest.[65] For each band, we define the count variable to only include originations of mortgages within the relevant LTV band. The estimates plotted in Figure C.3 show that there is no systematic correlation between LTV and evolution of the number of mortgages against flats. The contraction in flat originations around 2018Q4 appears approximately the same across all LTV bands. The lack of heterogeneity along LTV suggests that the central guidance did not change lenders' perceptions about borrower leverage. One explanation is that LTV might already reflect the effect of the shock as property prices dropped in responses to the shock. Lenders might have adjusted loan amounts in order to keep target LTV values constant.

[65] These are the same bands used by the Financial Conduct Authority (FCA) in their Quarterly Commentary on Mortgage lending statistics.

Bibliography

List of References

Almagro, M. and T. Dominguez-Iino. (2020). Location sorting and endogenous amenities: Evidence from amsterdam. Working Paper.

Amaral, F., M. Dohmen, S. Kohl, and M. Schularick (2025). Superstar Returns? Spatial Heterogeneity in Returns to Housing. Journal of Finance. Forthcoming.

Arnould, G., B. Guin, S. Ongena, and P. Siciliani (2020). (When) do banks react to anticipated capital reliefs? *Bank of England Staff Working Paper* (889).

Bakkensen, L. A. and L. Barrage (2021, 11). Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics. *The Review of Financial Studies 35*(8), 3666–3709.

Baldauf, M., L. Garlappi, and C. Yannelis (2020, 02). Does Climate Change Affect Real Estate Prices? Only If You Believe In It. *The Review of Financial Studies 33*(3), 1256–1295.

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: A structural model of the u.k. mortgage market. *The Journal of Finance 76*(6), 2997–3053.

Benetton, M., P. Bracke, J. F. Cocco, and N. Garbarino (2021, 11). Housing Consumption and Investment: Evidence from Shared Equity Mortgages. *The Review of Financial Studies 35*(8), 3525–3573.

Bernanke, B., M. Gertler, and S. Gilchrist (1996, February). The financial accelerator and the flight to quality. *Review of Economics and Statistics 78*(1), 1–15.

Bernstein, A., M. T. Gustafson, and R. Lewis (2019). Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics 134*(2), 253–272.

Binner, A. and B. Day (2015). Exploring mortgage interest deduction reforms: An equilibrium sorting model with endogenous tenure choice. *Journal of Public Economics 122*, 40–54.

Bishop, K., J. Dowling, N. V. Kuminoff, and A. Murphy (2023, November). Tax policy and the heterogeneous costs of homeownership. Working Paper 31824, National Bureau of Economic Research.

Blow, L. and L. Nesheim (2009). A retail price index including the shadow price of owner occupied housing. CEMMAP Working Paper CWP03, London.

Bracke, P. (2015). House prices and rents: Microevidence from a matched data set in central london. *Real Estate Economics 43*(2), 403–431.

Bracke, P. (2021). How Much Do Investors Pay for Houses? *Real Estate Economics 49*(S1), 41–73.

Bracke, P., E. W. Pinchbeck, and J. Wyatt (2018). The Time Value of Housing: Historical Evidence on Discount Rates. *The Economic Journal 128*(613), 1820–1843.

Bracke, P. and S. Tenreyro (2021, April). History Dependence in the Housing Market. *American Economic Journal: Macroeconomics 13*(2), 420–43.

Brunnermeier, M. K. (2009, March). Deciphering the liquidity and credit crunch 2007-2008. *Journal of Economic Perspectives 23*(1), 77–100.

Calder-Wang, S. (2020). The distributional impact of the sharing economy on the housing market. Working Paper.

Campello, M., R. A. Connolly, G. Kankanhalli, and E. Steiner (2022). Do real estate values boost corporate borrowing? Evidence from contract-level data. *Journal of Financial Economics 144*(2), 611–644.

Carney, M. (2015, 9). Breaking the tragedy of the horizon – climate change and financial stability. Speech by Mr Mark Carney, Governor of the Bank of England and Chairman of the Financial Stability Board, at Lloyd's of London, London, 29 September 2015.

Cerqueiro, G., S. Ongena, and K. Roszbach (2016). Collateralization, Bank Loan Rates, and Monitoring. *The Journal of Finance 71*(3), 1295–1322.

Chambers, M., C. Garriga, and D. E. Schlagenhauf (2009a). Accounting for changes in the homeownership rate. *International Economic Review 50*(3), 677–726.

Chambers, M. S., C. Garriga, and D. Schlagenhauf (2009b). The loan structure and housing tenure decisions in an equilibrium model of mortgage choice. *Review of Economic Dynamics*.

Chaney, T., D. Sraer, and D. Thesmar (2012, May). The Collateral Channel: How Real Estate Shocks Affect Corporate Investment. *American Economic Review 102*(6), 2381–2409.

Chavaz, M. (2016). Dis-integrating credit markets: diversification, securitization, and lending in a recovery. *Bank of England Staff Working Paper* (617).

Damen, S., M. Korevaar, and S. Van Nieuwerburgh (2025, February). An alpha in affordable housing? Working Paper 33470, National Bureau of Economic Research.

DeFusco, A. A., C. G. Nathanson, and E. Zwick (2022). Speculative dynamics of prices and volume. *Journal of Financial Economics 146*(1), 205–229.

Demers, A. and A. L. Eisfeldt (2022). Total returns to single-family rentals. *Real Estate Economics 50*(1), 7–32.

Diamond, R. and W. Diamond (2023). Racial Differences in the Total Rate of Return on Owner-Occupied Housing. Technical report.

Diamond, R. and W. Diamond (2024, May). Affordability of owner-occupied housing across us cities. *AEA Papers and Proceedings 114*, 340–45.

Diamond, R., T. McQuade, and F. Qian (2019, September). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review 109*(9), 3365–94.

DLUHC (2022, 7). Remediation costs: what leaseholders do and do not have to pay. Guidance Department for Levelling Up, Housing & Communities, published on 21 July 2022.

Díaz, A. and M. J. Luengo-Prado (2008). On the user cost and homeownership. *Review of Economic Dynamics 11*(3), 584–613.

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.

FCA (2022, 9). Report on insurance for multi-occupancy buildings. Report by the Financial Conduct Authority, published on 21 September 2022.

Ferri, G. and P. Murro (2015). Do firm–bank 'odd couples' exacerbate credit rationing? *Journal of Financial Intermediation 24*(2), 231–251.

Fu, Y. (1991). A model of housing tenure choice: Comment. *The American Economic Review 81*(1), 381–383.

Galiani, S., A. Murphy, and J. Pantano (2015). Estimating neighborhood choice models: Lessons from a housing assistance experiment. *The American Economic Review 105*(11), 3385–3415.

Garbarino, N. and B. Guin (2021). High water, no marks? Biased lending after extreme weather. *Journal of Financial Stability 54*, 100874.

Giglio, S., B. Kelly, and J. Stroebel (2021). Climate Finance. *Annual Review of Financial Economics 13*(1), 15–36.

Greenwald, D. and A. M. Guren (2024). Do credit conditions move house prices? Working Paper.

Gupta, A., H. Sapriza, and V. Yankov (2021). The collateral channel and bank credit. *Available at SSRN 4023809*.

Gyourko, J., C. Mayer, and T. Sinai (2013). Superstar cities. *American Economic Journal: Economic Policy 5*(4), 167–199.

Halket, J., L. Nesheim, and F. Oswald (2020). The housing stock, housing prices, and user costs: The roles of location, structure, and unobserved quality. *International Economic Review 61*(4), 1777–1814.

Halket, J. and M. Pignatti Morano di Custoza (2015). Homeownership and the scarcity of rentals. *Journal of Monetary Economics 76*, 107–123.

Han, L., R. Ngai, and K. Sheedy (2021). To Own or to Rent? The Effects of Transaction Taxes on Housing Markets. Working Paper.

Head, A., H. Lloyd-Ellis, and D. Stacey (2023). Heterogeneity, frictional assignment, and home-ownership. *International Economic Review 64*(3), 1265–1308.

Heckman, J. J., R. L. Matzkin, and L. Nesheim (2010). Nonparametric identification and estimation of nonadditive hedonic models. *Econometrica 78*(5), 1569–1591.

Henderson, J. V. and Y. M. Ioannides (1983). A model of housing tenure choice. *The American Economic Review 73*(1), 98–113.

Heston, A. and A. O. Nakamura (2011). Reported prices and rents of housing: Reflections of costs, amenities or both? In *Price and Productivity Measurement, Volume 1 - Housing*.

Higgins, B. (2023). Racial Segmentation in the US Housing Market. Technical report.

Hilber, C. A. (2005, March). Neighborhood externality risk and the homeownership status of properties. *Journal of Urban Economics 57*(2), 213–241.

Hilber, C. A. L. and A. Mense (2025). The rise and fall of the price-to-rent ratio: Why are superstar cities different? *Economic Journal*. Forthcoming.

Hilber, C. A. L. and W. Vermeulen (2016). The impact of supply constraints on house prices in england. *The Economic Journal 126*(591), 358–405.

Himmelberg, C., C. Mayer, and T. Sinai (2005, December). Assessing high house prices: Bubbles, fundamentals and misperceptions. *Journal of Economic Perspectives 19*(4), 67–92.

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.

Keys, B. J. and P. Mulder (2020, 10). Neglected No More: Housing Markets, Mortgage Lending, and Sea Level Rise. *NBER Working Paper 27930*.

Kirschenmann, K. (2016). Credit rationing in small firm-bank relationships. *Journal of Financial Intermediation 26*, 68–99.

Kiyotaki, N. and J. Moore (1997). Credit cycles. *Journal of Political Economy 105*(2), 211–248.

Kuchler, T. and B. Zafar (2019). Personal experiences and expectations about aggregate outcomes. *The Journal of Finance 74*(5), 2491–2542.

Kvaerner, J. S., N. Pavanini, and Y. Peng (2023). Leverage Regulation and Housing Inequality. Technical report.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015, April). The housing market(s) of san diego. *American Economic Review 105*(4), 1371–1407.

Levy, A. (2022). Housing policy with home-biased landlords: Evidence from french rental markets. Working Paper, June 2022.

McDonald, J. B. (1984). Some generalized functions for the size distribution of income. *Econometrica 52*(3), 647–663.

McDonald, J. B. and Y. J. Xu (1995). A generalization of the beta distribution with applications. *Journal of Econometrics 66*(1), 133–152.

Mian, A. and A. Sufi (2011). House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis. *American Economic Review 101*(5), 2132–2156.

Moktan, S., B. Guin, and L. Clarke (2025). The anatomy of a shock to residential real estate: the role of lending. Bank of England Staff Working Paper No. 1111.

Nguyen, D. D., S. Ongena, S. Qi, and V. Sila (2022, 03). Climate Change Risk and the Cost of Mortgage Credit. *Review of Finance 26*(6), 1509–1549.

Ortalo-Magné, F. and S. Rady (2006). Housing market dynamics: On the contribution of income shocks and credit constraints. *The Review of Economic Studies 73*(2), 459–485.

Ouazad, A. and M. E. Kahn (2021, 11). Mortgage Finance and Climate Change: Securitization Dynamics in the Aftermath of Natural Disasters. *The Review of Financial Studies 35*(8), 3617–3665.

Peydró, J.-L., F. Rodriguez-Tous, J. Tripathy, and A. Uluc (2023, 09). Macroprudential Policy, Mortgage Cycles, and Distributional Effects: Evidence from the United Kingdom. *The Review of Financial Studies*.

Poterba, J. and T. Sinai (2008). Tax expenditures for owner-occupied housing: Deductions for property taxes and mortgage interest and the exclusion of imputed rental income. *The American Economic Review 98*(2), 84–89.

Poterba, J. M. (1984, 11). Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach*. *The Quarterly Journal of Economics 99*(4), 729–752.

Poterba, J. M. (1992). Taxation and housing: Old questions, new answers. *The American Economic Review 82*(2), 237–242.

Scanlon, K. and C. Whitehead (2016, December). The profile of UK private landlords. Report.

Schüwer, U., C. Lambert, and F. Noth (2018, 04). How Do Banks React to Catastrophic Events? Evidence from Hurricane Katrina. *Review of Finance 23*(1), 75–116.

Sommer, K. and P. Sullivan (2018). Implications of us tax policy for house prices, rents, and homeownership. *American Economic Review*.

Sommer, K., P. Sullivan, and R. Verbrugge (2013). The equilibrium effect of fundamentals on house prices and rents. *Journal of Monetary Economics*.

Stiglitz, J. E. and A. Weiss (1981). Credit rationing in markets with imperfect information. *The American Economic Review 71*(3), 393–410.

Sufi, A. (2007). Information Asymmetry and Financing Arrangements: Evidence from Syndicated Loans. *The Journal of Finance 62*(2), 629–668.

Turner, T. M. (2003). Does investment risk affect the housing decisions of families? *Economic Inquiry 41*(4), 675–691.

Verbrugge, R. (2008). The puzzling divergence of rents and user costs, 1980–2004. *Review of Income and Wealth 54*(4), 671–699.

Verbrugge, R. and R. Poole (2010). Explaining the rent–oer inflation divergence, 1999–2007. *Real Estate Economics 38*(4), 633–657.