

NBER WORKING PAPER SERIES

LOCATION SORTING AND ENDOGENOUS AMENITIES:
EVIDENCE FROM AMSTERDAM

Milena Almagro
Tomás Domínguez-Iino

Working Paper 32304
<http://www.nber.org/papers/w32304>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024

First version: November 2019. We thank Guillaume Fréchet, Alessandro Lizzeri, Elena Manresa, Tobias Salz, Paul Scott, and Sharon Traiberman for their support at the early stages of this project. For useful feedback, we thank Juan Camilo Castillo, Don Davis, Fabian Eckert, Cecile Gaubert, Jessie Handbury, Allan Hsiao, Anders Humlum, Erik Hurst, Panle Jia-Barwick, Chad Syverson, and László Tétényi. Marek Bojko, Elias van Emmerick, Sara Gerstner and Sriram Tolety provided outstanding research assistance. We acknowledge financial support from the C.V. Starr Center for Applied Economics, the Liew Family Junior Faculty fellowship, and the George G. Rinder Faculty Fellowship. Any errors or omissions are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Milena Almagro and Tomás Domínguez-Iino. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Location Sorting and Endogenous Amenities: Evidence from Amsterdam
Milena Almagro and Tomás Domínguez-Iino
NBER Working Paper No. 32304
April 2024
JEL No. L0,L83,R0,R13,R20,R21

ABSTRACT

This paper shows the endogeneity of amenities plays a crucial role in determining the welfare distribution of a city's residents. We quantify this mechanism by building a dynamic model of residential choice with heterogeneous households, where consumption amenities are the equilibrium outcome of a market for non-tradables. We estimate our model using Dutch microdata and leveraging variation in Amsterdam's spatial distribution of tourists as a demand shifter, finding significant heterogeneity in residents' preferences over amenities and in the supply responses of amenities to changes in demand composition. This two-way heterogeneity dictates the degree of horizontal differentiation across neighborhoods, residential sorting, and inequality. Finally, we show the distributional effects of mass tourism depend on this heterogeneity: following rent increases due to growing tourist demand for housing, younger residents—whose amenity preferences are closest to tourists—are compensated by amenities tilting in their favor, while the losses of older residents are amplified.

Milena Almagro
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
Milena.almagro@chicagobooth.edu

Tomás Domínguez-Iino
5807 South Woodlawn Avenue
Booth School of Business
University of Chicago
Chicago, IL 60637
United States
tomasdi@uchicago.edu

1 Introduction

Socioeconomic inequality is tightly linked to residential choice, both across and within cities (Moretti, 2013). Higher socioeconomic status households can afford to live in locations with more desirable amenities. Furthermore, amenities improve as residential composition changes, reinforcing the desirability of locations. This response of a location’s amenities to demographic sorting has been shown to be a quantitatively important mechanism for amplifying welfare inequality (Guerrieri, Hartley and Hurst, 2013; Diamond, 2016). However, relatively little is understood about the nature of these *endogenous amenities*, as they are typically modeled as a one-dimensional object summarizing a wide variety of a location’s characteristics.

It is natural to think different types of households have diverse tastes for different types of consumption amenities, and that firms providing such amenities cater to this heterogeneity (George and Waldfogel, 2003). For example, when neighborhoods gentrify, the initial increase in the share of young, college-educated households is typically accompanied by an increase in the presence of bars and restaurants, and a reduction in mom-and-pop stores. While providing tractability, aggregating amenities into a single index does not allow for the *horizontal* differentiation of neighborhoods on the demand side, nor for differential supply-side responses to consumer heterogeneity. Moreover, if this heterogeneity plays an important distributive role, understanding its sources is crucial to design policies that alleviate urban inequality. For example, incumbent low-income residents living in gentrifying neighborhoods may not only suffer from higher housing prices, but also from the changes in neighborhood characteristics associated with the increase in higher-income households. Therefore, in this paper, we ask: How does preference heterogeneity over multiple endogenous consumption amenities shape within-city residential sorting and inequality?

To answer our research question, we build and estimate a dynamic spatial equilibrium model of a city with heterogeneity in household preferences over a *bundle of endogenous amenities*, whose supply caters to each neighborhood’s demographic composition. To estimate our model, we use restricted-access microdata from the Centraal Bureau voor de Statistiek (CBS), the statistics bureau of the Netherlands. From these data, we construct an annual panel of residential location choices for the universe of residents in the Netherlands. We complement these data with an annual panel of establishment counts, allowing us to track consumption amenities across time and space. Apart from the availability of high-quality data, Amsterdam provides an ideal laboratory to study the link between residential composition and endogenous amenities, as it has undergone significant changes due to the impact of mass tourism on local housing and amenity markets.

We start by showing the expansion of tourism across Amsterdam is significant enough to affect housing and local amenity markets. The number of overnight tourist stays went from 8 million in 2008 to nearly 16 million in 2017, along with a stark increase in housing units converted to short-term rentals (STR), primarily through the Airbnb platform. In contrast to hotels, which tend to spatially cluster in the city center, STR growth sprawled across all neighborhoods, reaching over 5% of the city-wide rental market and exceeding 20% in some central neighborhoods. Next, we

show STR expansion is sizable enough to impact rent prices. We continue by showing amenities catering to tourists increase in nearly every neighborhood, and their presence is negatively correlated with amenities catering exclusively to locals, such as nurseries/daycare facilities, which decline in more than half of neighborhoods at a median rate of -32%. Finally, we show different demographic groups respond differently to these neighborhood changes through their residential choices, suggesting different valuations for the changes in amenities.

In our model, residential choices and amenities are jointly determined equilibrium outcomes. We model the residential choices of local residents with a dynamic discrete choice setup—they are forward-looking, change locations subject to heterogeneous moving costs, and hold heterogeneous preferences over location attributes. We also specify a static model of how tourists choose the location where they book their STR. Hence, a location's total demand for housing and amenities is shaped by the location choices of both locals and tourists.

On the housing supply side, we assume atomistic absentee landlords supply housing to locals on traditional, long-term leases or to tourists on short-term leases. On the amenity supply side, monopolistically competitive firms provide a variety of consumption amenities that differentially cater to different types of locals and tourists. Compared to settings where amenities are collapsed to a one-dimensional quality index, introducing multiple types of amenities allows neighborhoods to endogenously become horizontally differentiated, because residents can trade off one type of amenity for another. This implies households of different income levels may disagree on which neighborhoods are most desirable, therefore decoupling income inequality from welfare (i.e, amenity-adjusted) inequality.

Because our micro-data tracks the residential locations of households, we can accommodate forward-looking behaviour and state-dependent moving costs in our estimation of locals' residential choices. These dynamic elements of our model are motivated by two features of our data. First, moving decisions are infrequent, suggesting significant moving costs. Second, we observe the probability of moving is state-dependent: it decreases in the time a household has been living in its current location. We capture these features of the data by i) including standard distance-adjusted moving costs, and ii) allowing agents to accumulate location capital that is lost upon moving, which introduces a dynamic, state-dependent component to moving costs. Failure to account for these dynamic elements is known to lead to biased estimates (Bayer, McMillan, Murphy and Timmins, 2016; Traiberman, 2019).

We estimate our dynamic location choice model by building upon the Euler Equation in Conditional Choice Probability (ECCP) methodology (Aguirregabiria and Magesan, 2013; Scott, 2013; Kalouptsi, Scott and Souza-Rodrigues, 2021b). We use an instrumental variable approach to address the endogeneity of rental prices and consumption amenities. Our demand estimates reveal preference parameters that correlate with demographics in reasonable ways. For example, households without children value restaurants the most, consistent with having the most leisure time among all groups. By contrast, households with children value nurseries the most. The highest

income and most educated households dislike touristic amenities.

On the amenity supply side, we also estimate reasonable supply responses of different amenity categories to different demographics. We find the presence of tourists mainly drives the entry of touristic amenities, restaurants, and non-food retail, but does not affect the entry of nurseries. Instead, the supply of nurseries responds most strongly to households with children, while younger households incentivize the entry of restaurants. The supply of grocery stores is the most homogeneous across household types, consistent with the notion they provide a service that is demanded similarly across socioeconomic strata.

We use our estimated model to run counterfactuals highlighting how preference heterogeneity and the endogeneity of amenities interact to determine sorting and inequality. In our first counterfactual, we compare the equilibrium outcome of our baseline specification with heterogeneous preferences to one with homogeneous preferences. We show that heterogeneous preferences lead to more spatial sorting, as households have more neighborhood dimensions along which to sort. However, although heterogeneous preferences and endogenous amenities can reinforce each other to generate more sorting, they can also reduce welfare inequality across household types. Intuitively, if preferences over amenities are misaligned between two demographic groups, then they sort into different locations. This sorting increases the supply of their most preferred amenities, making neighborhoods more differentiated, such that the two groups avoid competing for housing in the same location. Thus, there are two mechanisms reducing the welfare gap across groups when preferences are heterogeneous and amenities are endogenous: tailored amenities and lower rental prices. Our findings complement the existing literature on spatial sorting and inequality by introducing two-way heterogeneity in the relationship between households and amenities.

In our second counterfactual, we evaluate the effect of STR entry on local residents' welfare. We disentangle these effects into i) the direct effects on rent via the reduction in housing supply, and ii) the indirect effects on amenities via the endogenous response of amenity supply to the increased tourist population. The key insight behind our results is that while all residents lose from higher rents, their losses may be compensated or amplified depending on how they value the changes in amenities the tourists bring along. Moreover, we show the correlation between income and preferences for the amenities tourists bring determines how regressive STR entry is. If the lowest-income (highest-income) groups dislike the amenities that tourists bring, then STR entry is more regressive (progressive). Finally, in our third counterfactual we compare different forms of regulating mass tourism: through housing markets or amenity markets. Specifically, we compare a tax on short-term rentals to a tax on touristic amenities and show how the distributional impact of each policy lever depends on heterogeneity on both demand and supply sides of the amenities market.

Related literature. Spatial equilibrium models date back to [Rosen \(1979\)](#) and [Roback \(1982\)](#) and are a benchmark to study spatial inequality across and within cities ([Moretti, 2013](#); [Diamond, 2016](#); [Couture and Handbury, 2020](#)). A subset of the literature focuses on the within-city margin, but

typically remains silent on the exact mechanisms through which specific amenities are provided (Bayer, Ferreira and McMillan, 2007; Guerrieri et al., 2013; Ahlfeldt, Redding, Sturm and Wolf, 2015; Davis, Hartley, Gregory et al., 2019; Su, 2022). Recent studies impose structure on amenity provision, but often lack heterogeneity in residents' preferences over amenities or collapse amenities into a single quality index (Couture, Gaubert, Handbury and Hurst, 2021; Hoelzlein, 2020; Miyauchi, Nakajima and Redding, 2021). We contribute by allowing for preference heterogeneity over multiple and differentiated amenities, whose supply is microfounded through a market mechanism. We build upon the notion of "preference externalities": demand-side preference heterogeneity can translate into differences in the variety of products supplied (George and Waldfogel, 2003; Handbury, 2021). Similarly, we interpret neighborhoods as differentiated products where amenities play the role of endogenous product attributes, and highlight the implications for residential sorting and inequality.

Our paper also contributes to the literature on the STR industry, as well as tourism more broadly. There is extensive work on the effects of STR entry on the housing market (Sheppard, Udell et al., 2016; Koster, Van Ommeren and Volkhausen, 2021; Garcia-López, Jofre-Monseny, Martínez-Mazza and Segú, 2020; Barron, Kung and Proserpio, 2021) and hotel revenue (Zervas, Proserpio and Byers, 2017). Farronato and Fradkin (2018) study the effect of STR entry on competing hotel sector. Calder-Wang (2021) studies the distributional effects on the New York City rental market, focusing on rent effects but abstracting from amenity effects. Faber and Gaubert (2019) show the importance of tourism in the economic development of the Mexican coastline. Finally, Allen, Fuchs, Ganapati, Graziano, Madera and Montoriol-Garriga (2021) study the effects of seasonal tourism on prices of goods and amenities borne by residents of Barcelona. We complement their work by simultaneously studying the effects of tourism on both residential and amenity markets, showing how they interact to shape urban inequality.

In terms of methods, we use discrete-choice tools from the empirical industrial organization literature and show how they can be applied to urban residential markets (McFadden, 1974; Berry, 1994; Berry, Levinsohn and Pakes, 1995; Rust, 1987). Specifically, our dynamic estimation uses the Euler Equation in Conditional Choice Probabilities (ECCP) estimator (Hotz and Miller, 1993; Arcidiacono and Miller, 2011; Aguirregabiria and Magesan, 2013; Scott, 2013; Kalouptsi et al., 2021b). The method has been applied to several contexts where dynamics are first order: agricultural markets (Scott, 2013; Hsiao, 2021), occupational choice (Traiberman, 2019; Humlum, 2021), and residential choice (Diamond, McQuade and Qian, 2019; Davis et al., 2019; Davis, Gregory, Hartley and Tan, 2021).

2 Data

Individual-level data: residential histories and socioeconomic characteristics. Our individual-level microdata is from the statistical bureau of the Netherlands, Centraal Bureau voor de Statistiek (CBS). The key dataset for our dynamic model is the residential cadaster, from which we construct

a panel of residential history for the universe of individuals in the Netherlands. We also observe household-level demographics from tax return data: income, educational attainment, employment status, household composition, and ethnic background. We classify households as low-, medium-, or high-skill using educational attainment bins. Because we do not observe workplace nor occupations, our analysis focuses on residential market, rather than labor market outcomes. Further details are in Appendix [A.2.1](#).

Housing unit data: tax valuations, tenancy status, physical characteristics, rental prices, and transaction values. First, we obtain property values from a CBS tax appraisal panel for the universe of residential housing units for 2006-2020, which also includes geo-coordinates, quality measures, and the occupant’s tenancy status (owner-occupied, rental, social housing). For the subset of these properties that are transacted, we can confirm that their tax appraisals are highly correlated with transaction prices (we observe all housing sale transactions in the Netherlands). Second, we obtain rental prices from a CBS national rent survey for 2006-2019. Since the survey does not cover the universe of tenants, we impute rental prices by linking it to the universe of tax appraisal valuations and employing a random forest, which outperforms traditional linear hedonic models ([Mullainathan and Spiess, 2017](#)). Imputation details are in Appendix [A.2.4](#).

Neighborhood-level data: amenities, demographic changes, tourist inflows. We use two levels of geographic units based on Amsterdam’s administrative divisions: 99 “wijk” (neighborhoods) that belong to 25 larger “gebied” (districts). An average wijk had roughly 8,540 inhabitants as of 2018. After dropping unpopulated or industrial-use-only neighborhoods, we end up with 95 neighborhoods and 22 districts. We observe annual neighborhood-level outcomes from Amsterdam City Data (ACD) from 2008-2018. These outcomes include demographics (e.g., ethnic, income, and skill composition) and a rich set of consumption amenities. We also obtain city-level tourist inflows from ACD. The ACD wijk-level and Tourism data are publicly available at [ACD BBGA](#) and [ACD Tourism](#).

For our estimation procedure and counterfactuals, we narrow down the set of amenities to six: restaurants, bars, food stores, non-food stores, nurseries, and “touristic amenities”. First, we chose these categories because they are available at a granular spatial unit for the whole time period in our sample (many categories are not reported every year nor at every administrative subdivision). Second, these categories likely vary in the extent to which they cater to tourists versus different types of locals. “Touristic amenities” is a category defined by ACD that includes tourist-oriented business such as travel agencies, cultural/recreational establishments, and lodging. We remove lodging from the original ACD definition because we treat hotels separately in our analysis—we consider them solely as accommodation for tourists rather than as a consumption amenity that could potentially be valued by both tourists and locals. Thus, our final measure of touristic amenities consists of consumption services that some locals may value, such as cultural/recreational establishments. Bars includes pub-style establishments that serve only alcohol, as well as cafe-style establishments that serve both coffee and alcoholic drinks, without being full-fledged restaurants.

Food stores refers to establishments that sell food without service, such as a grocery or convenience store. Non-food stores refers to non-food commercial retail, such as clothing stores. Restaurants and nurseries are self-explanatory.

Short-term rental listings. Airbnb holds over 80% of the STR market share in Amsterdam. Hence, throughout the paper we use Airbnb and STR interchangeably. Our Airbnb data is from [Inside Airbnb](#), an independent website providing monthly web-scraped listings data for many cities. Our data consist of listing-level observations with information such as geo-coordinates, prices per night, calendar availability, and reviews. We use this information to separately identify “active” from “dormant” STR listings, and to flag commercially-operated listings—those likely to be permanently rented to tourists, thus reducing housing supply for locals. We define commercial listings as entire-home listings with booking activity above a threshold. Classification details are in Appendix [A.2.7](#).


Final sample: time period and geographic unit of analysis. We construct an annual panel of location choices and characteristics for 2008-2018. For our dynamic model of local’s residential choices, we aggregate 95 neighborhoods (wijk) into 22 districts (gebied). Using larger geographical units allows us to estimate more precise *conditional choice probabilities* that feed into the estimation of the dynamic model. Because demand is at the district level, we also use districts in the estimation of amenity supply. Our estimation of housing supply and tourist demand only requires *unconditional choice probabilities*. Thus, for those cases we use the smaller neighborhoods as spatial units, allowing us to obtain more precise estimates.

3 Stylized facts

We present the stylized facts of our empirical setting and how they motivate our model’s key features. We show tourism volume and STR penetration have grown over time and across neighborhoods, and how such growth correlates with our outcomes of interest: rental prices, consumption amenities, and the socioeconomic composition of residents. The role these tourism-induced compositional changes have in shaping local amenities, and how local residents respond to such amenity changes by moving, is what motivates our overarching question of how endogenous amenities interact with sorting across neighborhoods.

Fact 1: Tourists and STR listings have grown dramatically and sprawled across Amsterdam. Amsterdam has one of the highest tourist-to-local ratios in the world, above Florence and slightly below Venice (source: [ESTA](#)). Figure 1 shows that, between 2008-2017, the number of overnight stays per resident doubled, hotel capacity grew from approximately 22,000 to 31,000 rooms, while STR listings grew from zero to over 25,000. The figure also shows the evolution of commercially-operated listings, which are available year-round and therefore comparable to hotel rooms in terms of nights-availability. By 2017, there were approximately 7,000 of these listings, which is equivalent to 25% of the city’s stock of hotel rooms.

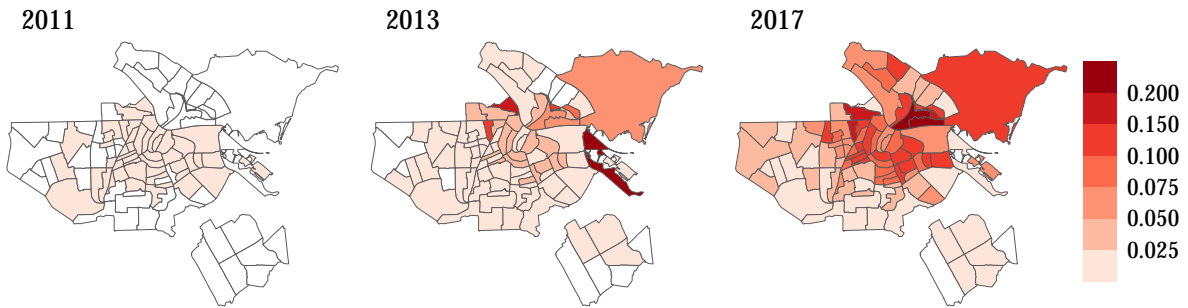
Figure 1: Overnight stays per resident, hotel rooms, and STR listings (2008-2017).



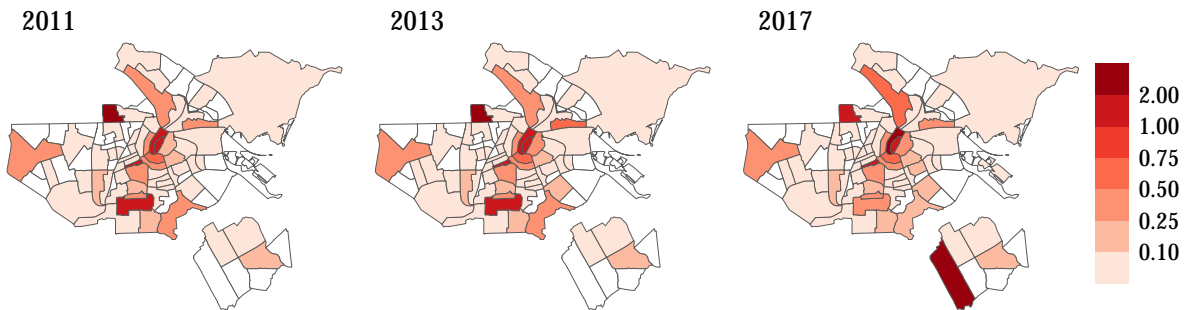
Notes: On the left, “overnight stays per 100 residents” is constructed as annual overnight stays (in hotels and STR) divided by population, and multiplied by 100—a value of 5 means that on an average night there are 5 tourists per 100 residents. On the right, active and commercial Airbnb listings are constructed from Inside Airbnb data using the procedure described in Appendix A.2.7. Hotel, stay and population data are from [ACD Tourism](#) and [ACD BBGA](#).

Figure 2: STR share of rental stock and hotel beds per resident (2011-2017).

Panel A: Commercial STR share of rental stock (2011-2017).



Panel B: Hotel beds per resident (2011-2017).



Notes: Maps show neighborhood (“wijk”) level outcomes. We construct commercial STR listings from Inside Airbnb data, using the procedure described in Appendix A.2.7. Rental housing stock, hotel beds, and population data is from [ACD BBGA](#). The rental stock corresponds to private market rental stock (i.e., we exclude social housing rentals).

Figure 2 shows commercial listings have sprawled to cover most of the city. By contrast, the spatial distribution of hotels remains mostly unchanged and clustered in the city center. This is partly due to zoning regulations that apply to hotels but not to the STR segment. At the aggregate level, commercially-operated STR listings represented 6% of the rental market in 2017, exceeding 20% in some central neighborhoods. These trends suggest the increasing presence of tourists as part of the city’s population is significant enough to alter local housing and amenity markets.

Fact 2: Rents have increased more in neighborhoods with more STR entry. Table 1 shows the intensity of STR penetration is positively correlated with housing market outcomes. OLS regressions in the top panel show a 1% increase in a neighborhood’s commercial STR listings is associated with a rent increase between .06-.11%. These magnitudes are sizable given rents grew at an annualized rate of 1.02% between 2009-2019, and are also in line with recent studies estimating the effect of STR on housing market prices. For example, [Barron et al. \(2021\)](#) estimate an STR elasticity of rent of 0.018. The bottom panel of Table 1 repeats the regression exercise for sale prices, finding a 1% increase in commercial STR listings is associated with a house sale price increase between .04-.11% in OLS specifications.

Table 1: Airbnb intensity and housing market outcomes

	Ln (rent/m2)					
	OLS	IV	OLS	IV	OLS	IV
Ln (commercial Airbnb listings)	0.065*** (0.008)	0.091*** (0.021)	0.051*** (0.006)	0.114*** (0.021)	0.109*** (0.018)	0.205* (0.093)
Control variables			X	X	X	X
District-year FE					X	X
First stage F-stat		586.89		384.21		69.66
Observations	770	770	763	763	763	763

	Ln (house sale price)					
	OLS	IV	OLS	IV	OLS	IV
Ln (commercial Airbnb listings)	0.109*** (0.016)	0.290*** (0.030)	0.034*** (0.006)	0.149*** (0.016)	0.037* (0.022)	0.326** (0.102)
Control variables			X	X	X	X
District-year FE					X	X
First stage F-stat		572.02		370.87		65.9
Observations	738	738	737	737	737	737

Notes: Observations are at the wijk (neighborhood) level. A “district” is a larger spatial unit than a neighborhood. Rent prices are neighborhood-average long-term rental prices constructed from CBS rent surveys. House sale prices are neighborhood average transaction values, constructed from CBS data covering the universe of housing transactions. Commercial Airbnb listings are constructed from the Inside Airbnb data (see Appendix A.2.7 for construction details). Neighborhood-level control variables are: housing stock, average income, high-skill population share, all from [ACD BBGA](#). Standard errors are clustered at the wijk level in parenthesis.

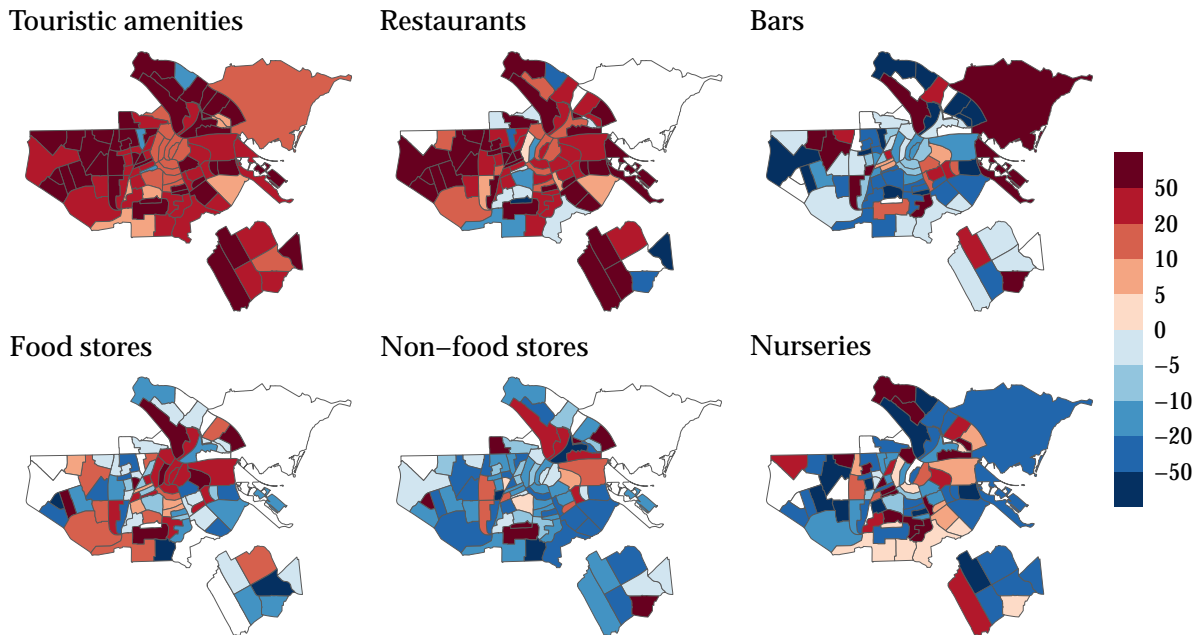
The main endogeneity concern from the OLS results is that time-varying neighborhood-level unobservables correlate with both STR penetration and housing market prices, leading to biases that depend on the sign of such correlations. For example, if neighborhoods that are becoming

more attractive to tourists are becoming less attractive to locals, then such areas will have more STR and lower rent, leading to downward-biased OLS estimates. Beyond including controls that likely correlate with such unobservables, we address these concerns with a shift-share instrument, a common research design in the literature measuring the impact of STR on housing markets (Barron et al., 2021; Garcia-López et al., 2020).

The “shift” part of the instrument exploits time variation in worldwide demand for STR, as proxied by online search activity for Airbnb. The “share” part constructs neighborhood-level exposure to tourism by using the spatial distribution of historic monuments. Our exclusion restriction requires both factors to be orthogonal to time-varying and neighborhood-level unobservables, conditional on the rest of the covariates. First, Airbnb’s worldwide popularity is unlikely to be informative of neighborhood-specific trends. Second, the spatial distribution of monuments determined centuries ago is unlikely to be informative of current trends affecting housing prices. Our results indicate the OLS estimates are downward-biased. This is consistent with the unobservables being positively correlated with Airbnb presence and negatively correlated with housing market prices, i.e., they are likely dis-amenities for local residents.

Finally, note the reduced-form results from Table 1 capture the total impact of STR. This is a combination of i) less housing supply for locals, which raises rents, and ii) changes in amenities, which can raise or lower rents depending on how locals value such amenities. This limitation of the reduced-form analysis is what motivates our model, with which we aim to disentangle these two channels.

Figure 3: Evolution of consumption amenities (2011-2017 pp changes).

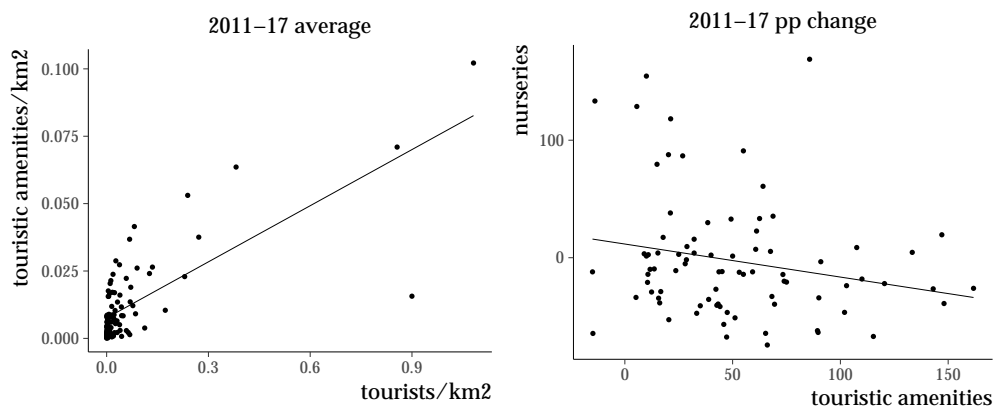


Notes: Maps show percentage point changes between 2011-2017 for each amenity sector. Data is from ACD BPGA.

Fact 3: Amenities have tilted towards tourists and away from locals. Beyond the impact of STR on the housing market, the amenities surrounding the housing units have also changed as tourists become an increasing share of the city’s population.

Figure 3 shows touristic amenities have grown across nearly all neighborhoods, although at different intensities, while amenities catering exclusively to locals, such as nurseries, have declined in most locations. Figure 4 confirms touristic amenities indeed locate in neighborhoods with high tourist intensity, and that their growth is negatively correlated with amenities that are clearly targeted to locals, such as nurseries. Overall, these patterns are consistent with tourists having different preferences over amenities than locals. As for the other 4 amenities displayed in Figure 3, they are likely in between the two extremes of touristic amenities and nurseries in the sense they would not a-priori seem to cater solely to locals or solely to tourists, but likely to both. The purpose of our amenity supply model is to estimate the extent to which amenities such as these lie in between the two extremes.

Figure 4: Spatial correlation between tourist-oriented and local-oriented amenities.

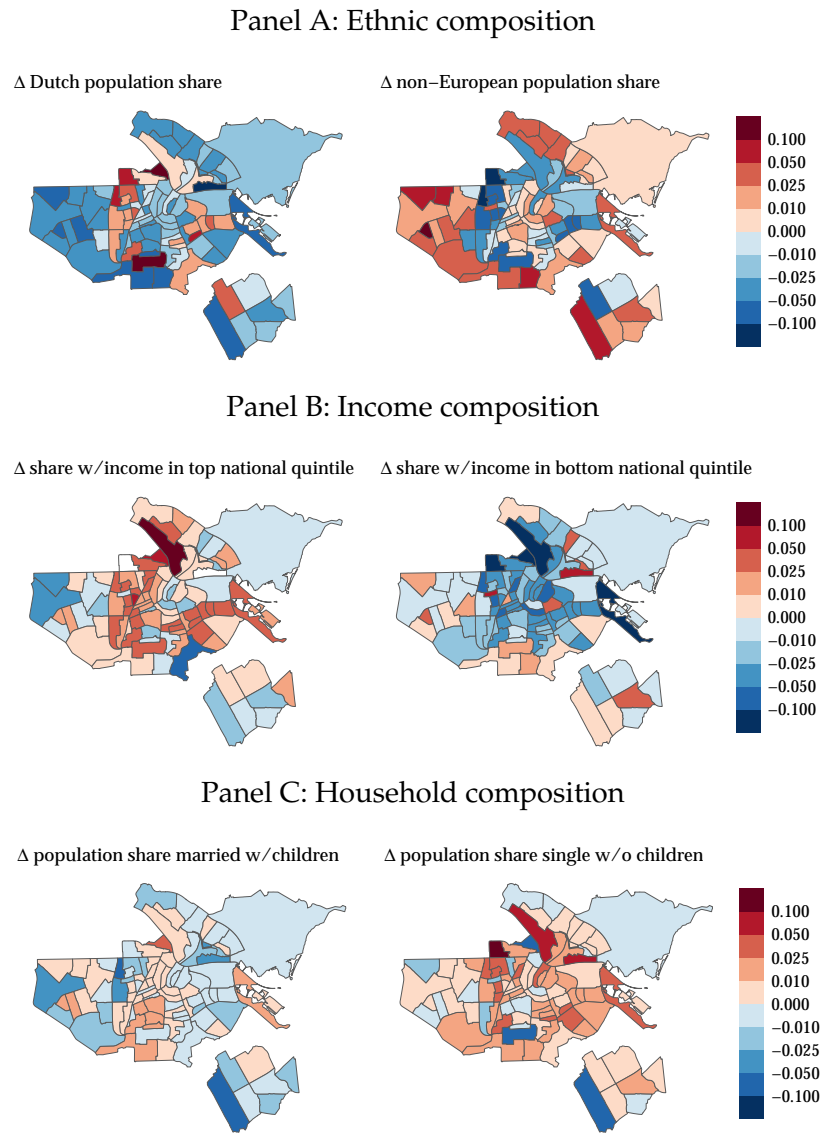


Notes: Left figure plots the 2011-2017 average tourist intensity vs touristic amenity intensity, for each neighborhood. Right figure plots the 2011-2017 percentage point change for nurseries vs touristic amenities, for each neighborhood. Nurseries decline in 58% of neighborhoods, with a median decline of -32%. Data is from [ACD BBGA](#).

Fact 4: The composition of residents has changed heterogeneously across neighborhoods. Figure 5 shows how socioeconomic composition has evolved in each neighborhood by displaying changes in the population shares of various demographic groups. The top panel shows a falling share of residents with Dutch background in most neighborhoods, except those around the city center. By contrast, the share of non-European immigrants has increased in a few central neighborhoods and mostly in the periphery. In terms of income heterogeneity, the middle panel shows the share of residents in the top 20% of the national income distribution has grown in central neighborhoods but not in the outskirts, indicating a rise in income inequality between the core and periphery. The bottom panel shows heterogeneity along household composition: households

with children have become increasingly outnumbered by those without in most neighborhoods.¹ To summarize, the heterogeneity in the socioeconomic make-up of neighborhoods and in their evolution over time motivates the heterogeneity in our model’s demand primitives: rent elasticities, moving costs, and valuation of amenities.

Figure 5: Changes in socioeconomic composition of neighborhoods (2011-2017).



Notes: Maps shows changes in neighborhood population share of each group. Data is from [ACD BBGA](#).

¹It is worth noting that changes in neighborhood composition can occur because households move, but also because household characteristics can change for those that do not move. For fixed characteristics such as ethnicity we can guarantee all the compositional change is due to moving, but this won’t be the case for mutable characteristics such as income or marital status. From the aggregate data with which Figure 5 is constructed, we cannot disentangle how much of the compositional changes along mutable characteristics comes from households moving versus their status changing. However, we can isolate the moving component in the estimation of our structural model by leveraging individual-level data that explicitly tracks residential location over time.

4 A dynamic model of an urban rental market

Motivated by the previous facts, we build a dynamic model of a city’s rental market that consists of i) heterogeneous households and tourists making location decisions, ii) landlords who can rent their units to locals or tourists, and iii) a market for amenities that microfounds how the composition of amenities endogenously responds to the composition of locals and tourists.

Notation. There are $J + 1$ locations: J locations inside the city and an outside option. Households are classified into $K + 1$ types: K different types of locals and a tourist type T , each differing in their preferences over consumption amenities. We define the population composition of location j at time t as the following vector,

$$M_{jt} \equiv [M_{jt}^1, \dots, M_{jt}^K, M_{jt}^T]', \quad (1)$$

where M_{jt}^k is the number of type k households in location j .

Consumption amenities are classified into S sectors, each consisting of multiple firms providing their own differentiated varieties. For example, if the sector is “restaurants”, a firm is an individual restaurant supplying its own variety. Hence, we use the terms “firm” and “variety” interchangeably. Let $N_{s jt}$ denote the number of varieties in sector s and location j at time t . We define the amenities of location j as the vector that lists the number of varieties in each sector,

$$a_{jt} \equiv [N_{1 jt}, \dots, N_{S jt}]'. \quad (2)$$

We present the model in three steps. First, section 4.1 shows how amenities a_{jt} are endogenously determined by the population composition M_{jt} . Second, sections 4.2-4.3 show the reverse mapping—how population composition adjusts to amenities through location choices. Third, section 4.4 brings the two mappings together by providing an equilibrium definition through which population composition and amenities are jointly determined.

4.1 Endogenous amenities

This section shows how amenities are endogenously determined by residential composition. We present main results, relegating derivations to Appendix A.3.1.

Demand for amenities. Households have Cobb-Douglas preferences over housing H and a composite of consumption amenities C , with ϕ^k being the expenditure share on C for a type k household. Let w_t^k denote the type k income at time t , so that total expenditures on housing and the amenities composite are $(1 - \phi^k)w_t^k$ and $\phi^k w_t^k$, respectively. Next, conditional on picking location j , a type k consumer chooses how much of her after-rent income $\phi^k w_t^k$ to allocate across the locally available amenity sectors and varieties. We assume consumers hold *Cobb-Douglas preferences across amenity sectors* and *CES preferences over varieties within a sector*. Hence, the quantity demanded by

type k for variety i in sector-location sj at time t is,

$$q_{isjt}^k = \frac{\alpha_s^k \phi^k \omega_t^k}{p_{isjt}} \left(\frac{p_{isjt}}{P_{sjt}} \right)^{1-\sigma_s}, \text{ with } P_{sjt} \equiv \left(\sum_{i=1}^{N_{sjt}} p_{isjt}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}}, \quad (3)$$

where α_s^k is the sector's budget share, $\sigma_s > 1$ is the substitution elasticity across varieties within the sector, p_{isjt} is the variety price, N_{sjt} is the number of varieties in sector-location sj , and P_{sjt} is the sector-location price index. Total demand for variety i is obtained by scaling up individual demand 3 by the location's type k population and aggregating across types,

$$q_{isjt} = \sum_k q_{isjt}^k M_{jt}^k. \quad (4)$$

Supply of amenities. Firms are indexed by i and engage in monopolistic competition with free entry, facing the same marginal cost c_{sjt} within a sector-location sj . This implies every firm i in sector-location sj chooses the same price and quantity,²

$$p_{isjt} = \frac{c_{sjt}}{1 - \frac{1}{\sigma_s}} \quad \forall i \in sjt \implies p_{isjt} = p_{sjt} \quad \text{and} \quad q_{isjt} = q_{sjt} \quad \forall i \in sjt. \quad (5)$$

Apart from variable costs, firms also pay a fixed cost per period F_{sjt} which captures operational costs such as the cost of renting commercial space. Under these assumptions, firms enter until the following zero-profit condition holds,

$$(p_{sjt} - c_{sjt})q_{sjt} = F_{sjt}(N_{jt}), \quad \text{where } N_{jt} = \sum_s N_{sjt}. \quad (6)$$

Note we assume F_{sjt} is increasing in the location's total number of firms across all sectors, N_{jt} . This allows for congestion forces, such as competition for commercial real estate among firms operating in the same location. Alternatively, congestion could be sector-specific (a function of N_{sjt}). We opt for our baseline assumption given we expect all firms in a location to compete in the same real estate market.

Equilibrium amenities. The market clearing conditions for the amenities market are obtained by plugging 4 and 5 in 6. This delivers the equilibrium number of varieties N_{sjt} as a function of population composition M_{jt} ,

$$N_{sjt} = \frac{1}{\sigma_s F_{sjt}} \sum_k \alpha_s^k \phi^k \omega_t^k M_{jt}^k. \quad (7)$$

²Our amenity data do not contain the firm-level data required to accommodate within sector-location price differences. Given the data limitations, our assumption on marginal costs allows for an empirically tractable mapping of how amenities respond to demographic composition that still allows for heterogeneity in prices across sector-location-time.

Given our definition for amenities in 2, this section has constructed a mapping, which we denote $\mathcal{A}(\cdot)$, that goes from population composition M_{jt} to amenities a_{jt} by imposing market clearing in the amenities market,

$$a_{jt} = \mathcal{A}(M_{jt}). \quad (8)$$

4.2 Housing supply

Let \mathcal{H}_{jt} denote the total housing stock, measured in units of floor space, in location j and year t . We assume housing stock is inelastic in the short-run and follows an exogenously determined path over time.³

In each location there is a continuum of absentee landlords making a binary choice between renting in the long-term market (LT) or in the short-term market (ST). The income obtained per unit of floor space from long-term rentals is r_{jt} , and from short-term rentals is p_{jt} . Given different matching and managerial costs involved in renting short- versus long-term, we introduce a wedge in operating costs κ_{jt} between the two choices. An individual landlord's problem is therefore,

$$\max \{ \alpha r_{jt} + \epsilon_{LT}, \quad \alpha p_{jt} - \kappa_{jt} + \epsilon_{ST} \},$$

where α is the marginal utility of income and ϵ_{LT} and ϵ_{ST} are idiosyncratic shocks.

Housing supply in each location. Under the assumption that the idiosyncratic shocks are distributed type I EV, the amount of housing supplied (in units of floor space) in the long- and short-term markets are, respectively,

$$\mathcal{H}_{jt}^{LT,S}(r_{jt}, p_{jt}) = \frac{\exp(\alpha r_{jt})}{\exp(\alpha r_{jt}) + \exp(\alpha p_{jt} - \kappa_{jt})} \mathcal{H}_{jt}, \quad (9)$$

$$\mathcal{H}_{jt}^{ST,S}(r_{jt}, p_{jt}) = \mathcal{H}_{jt} - \mathcal{H}_{jt}^{LT,S}(r_{jt}, p_{jt}). \quad (10)$$

4.3 Housing demand

Demand for housing is composed of the demand from local residents and the demand from tourists staying in short-term rental units.

4.3.1 Demand from locals

At the beginning of each period t , a household i chooses a residential location j_{it} . Locations inside the city are indexed $j = 1, \dots, J$, while the outside option of leaving the city is denoted $j = 0$.

Moving costs and location tenure. Upon moving, households incur a moving cost that consists of

³On average, annual growth of housing stock in Amsterdam is 1.2%, similar to the 0.9% value for San Francisco, one of the least housing-elastic cities in the US (sources: datacommons.org, [Building Permits Survey](#), [Saiz \(2010\)](#)). Our assumption of inelastic housing supply is broadly in line with other studies of housing supply in the Netherlands ([Vermeulen and Rouwendal, 2007](#)).

a fixed component and a bilateral distance-adjusted component,

$$MC^k(j_{it}, j_{it-1}) = \begin{cases} 0 & \text{if } j_{it} = j_{it-1} \\ m_0^k + m_1^k \text{dist}(j_{it}, j_{it-1}) & \text{if } j_{it} \neq j_{it-1} \text{ and } j_{it}, j_{it-1} \neq 0 \\ m_2^k & \text{if } j_{it} \neq j_{it-1}, \text{ and } j_{it} = 0 \text{ or } j_{it-1} = 0, \end{cases}$$

where $\text{dist}(j_{it}, j_{it-1})$ is distance between current and previous location, and m_0^k and m_2^k are fixed costs of moving within and outside the city, respectively. Moreover, households accumulate tenure by staying in a location. Tenure is key to rationalize the decreasing hazard rate of moving in the data, and evolves deterministically as,

$$\tau_{it} = \begin{cases} \min\{\tau_{it-1} + 1, \bar{\tau}\} & \text{if } j_{it} = j_{it-1} \\ 1 & \text{otherwise.} \end{cases}$$

We assume tenure can be accumulated up to a maximum absorbing state $\bar{\tau}$. Note $MC^k(j_{it}, j_{it-1})$ can be interpreted as the static component of the cost of moving, as in static migration models (Bryan and Morten, 2019), while τ_{it} is the dynamic component. That is, for a household that remains in the same location over multiple periods the term $MC^k(j_{it}, j_{it-1})$ remains constant over time while τ_{it} evolves.

Individual state variables. We denote $x_{it} \equiv (j_{it-1}, \tau_{it-1})$ as the vector of individual state variables that are observable to the econometrician: location j_{it-1} and tenure τ_{it-1} . Households also face a vector of unobservable idiosyncratic preference shocks for each location, $\epsilon_{it} = [\epsilon_{i0t}, \epsilon_{i1t}, \dots, \epsilon_{ijt}]$.

Aggregate state variables. We denote $\omega_t \equiv (r_t, a_t, b_t, \xi_t)$ as the matrix of aggregate state variables, which consists of the vector of rent $r_t = (r_{1t}, \dots, r_{jt})$, the matrix of consumption amenities $a_t = [a_{1t}, \dots, a_{jt}]$, exogenous location attributes $b_t = [b_{1t}, \dots, b_{jt}]$, and factors that are unobservable to the econometrician $\xi_t = [\xi_{1t}, \dots, \xi_{jt}]$. In what follows and to condense notation, we denote with subscript t the functions that depend on ω_t , in particular the flow utility and value function,

$$u_t^k(j, x_{it}) \equiv u^k(j, x_{it}, \omega_t) \quad \text{and} \quad V_t^k(x_{it}, \epsilon_{it}) \equiv V^k(x_{it}, \epsilon_{it}, \omega_t).$$

Flow utility and value function. The flow payoff of a household i of type- k in location j is a function of its individual state x_{it} and the aggregate state at time t ,

$$u_t^k(j, x_{it}) = \bar{u}_t^k(j) + \delta_\tau^k \log \tau_{it} - MC^k(j, j_{it-1}), \quad (11)$$

where we define $\bar{u}_t^k(j)$ as the component of payoffs that is common to all type k households. We stress that aggregate state variables such as r_{jt} , a_{jt} , and b_{jt} enter flow payoffs through $\bar{u}_t^k(j)$. Also note the k index in $\bar{u}_t^k(j)$ implies preferences heterogeneity over such variables. We allow for the utility of households to increase with location capital, motivated by the fact that the likelihood

of moving to a new location decreases with location tenure (see Appendix A.2.5 for evidence). Finally, household i 's dynamic problem can be written recursively as,

$$V_t^k(x_{it}, \epsilon_{it}) = \max_{j \in \{0, 1, \dots, J\}} u_t^k(j, x_{it}) + \epsilon_{ijt} + \beta \mathbb{E}_t \left[V_{t+1}^k(x_{it+1}, \epsilon_{it+1}) | j, x_{it}, \epsilon_{it} \right].$$

Location choice and evolution of the population. If $\epsilon_{ijt} \stackrel{i.i.d.}{\sim}$ type I EV, the probability a type k household chooses j , conditional on state x_{it} , is,

$$\mathbb{P}_t^k(j | x_{it}) = \frac{\exp \left(u_t^k(j, x_{it}) + \beta \mathbb{E}_t \left[V_{t+1}^k(x_{it+1}, \epsilon_{it+1}) | j, x_{it}, \epsilon_{it} \right] \right)}{\sum_{j'} \exp \left(u_t^k(j', x_{it}) + \beta \mathbb{E}_t \left[V_{t+1}^k(x_{it+1}, \epsilon_{it+1}) | j', x_{it}, \epsilon_{it} \right] \right)}. \quad (12)$$

To keep track of the population distribution over time, let $\pi_t^k(j, \tau)$ denote type k 's joint probability of living in location j with tenure τ at the end of period t . We can write how this object evolves by using the conditional choice probability 12,

$$\pi_t^k(j, \tau) = \begin{cases} \sum_{\tau'} \sum_{j' \neq j} \mathbb{P}_t^k(j | j', \tau') \pi_{t-1}^k(j', \tau') & \tau = 1 \\ \mathbb{P}_t^k(j | j, \tau - 1) \pi_{t-1}^k(j, \tau - 1) & \tau \in [2, \bar{\tau}) \\ \mathbb{P}_t^k(j | j, \bar{\tau} - 1) \pi_{t-1}^k(j, \bar{\tau} - 1) + \mathbb{P}_t^k(j | j, \bar{\tau}) \pi_{t-1}^k(j, \bar{\tau}) & \tau = \bar{\tau}. \end{cases} \quad (13)$$

Finally, to obtain the type k population count for location j we scale probabilities in 13 by M_t^k , the total number of type k locals city-wide, and sum across tenure states,

$$M_{jt}^k(r_t, a_t) = \sum_{\tau} \pi_t^k(j, \tau) M_t^k \quad \forall k \in \{1, \dots, K\}. \quad (14)$$

The left-hand side of the equation above is explicit on location choices depending on the distribution of rent and amenities, given the $\pi_t^k(j, \tau)$ term on the right-hand side depends on the choice probability 12, which in turn depends on r_t and a_t .

Housing demand from locals in each location. Note equation 14 is the demand for location j measured in units of households, not in units of floor space. Hence, we need to define the floor space demanded by each type of household. Recall from section 4.1 that households have Cobb-Douglas preferences over housing and amenity consumption, which implies a type- k household in location j consumes $f_{jt}^k \equiv \frac{(1-\phi^k)w_t^k}{r_{jt}}$ units of floor space. Therefore, long-term rental demand from locals for location j , measured in units of floor space, is,

$$\mathcal{H}_{jt}^{LT,D}(r_t, a_t) = \sum_{k=1}^K M_{jt}^k(r_t, a_t) f_{jt}^k. \quad (15)$$

4.3.2 Demand from tourists

There is an exogenous number of tourists M_t^T arriving into the city and choosing to stay in a short-term rental or a hotel.

Tourists in short-term rentals. Tourists staying in a STR in location j obtain the following payoff (excluding idiosyncratic shocks),

$$u_{jt}^{ST} = \delta_j^{ST} + \delta_t^{ST} + \delta_p^{ST} \log p_{jt} + \delta_a^{ST} \log a_{jt} + \zeta_{jt}^{ST}, \quad (16)$$

where p_{jt} is the location's short-term rental prices, a_{jt} are amenities, and the remaining terms are factors that are unobservable to the econometrician, which we incorporate with fixed effects $(\delta_j^{ST}, \delta_t^{ST})$ and a time-varying location quality ζ_{jt}^{ST} . The payoff in 16 is also subject to a type I EV idiosyncratic shock ε_{jt}^{ST} , which gives a closed-form expression of the number of tourists choosing to stay in a STR in j ,

$$M_{jt}^{ST}(p_t, a_t) = \frac{\exp(u_{jt}^{ST})}{\sum_{j'=0}^J \exp(u_{j't}^{ST})} \cdot M_t^T. \quad (17)$$

It is through equation 17 that the spatial distribution of tourists in short-term rentals responds to changes in short-term rental prices p_t and amenities a_t .

Tourists in hotels. Tourists also have the option to stay in a city-wide hotel sector, which we treat as an outside option. While this endogenizes the city-wide number of tourists in hotels, it does not endogenize how they are distributed across locations. We distribute tourists in hotels across locations in proportion to the hotel capacity observed in the data. We take this approach because we do not have hotel price data nor bookings to estimate hotel demand across locations. Although city-wide hotel capacity increases during our sample period, we consider our approach a reasonable solution given the spatial distribution of hotels does not substantially change and most of the spatial expansion of tourist accommodation occurs through short-term rentals (our stylized fact 1 from section 3). Operationally, we denote the hotel option as an outside option with its payoff normalized to zero. Hence, the number of tourists who endogenously choose the hotel sector at the aggregate city level is the residual of those choosing short-term rentals,

$$M_t^H(p_t, a_t) = M_t^T - \sum_{j=1}^J M_{jt}^{ST}(p_t, a_t).$$

The tourist population in hotels in location j is constructed as $M_{jt}^H(p_t, a_t) = s_{jt}^{beds} \times M_t^H(p_t, a_t)$, where s_{jt}^{beds} is the location j share of the city's hotel beds observed in the data. Finally, we obtain the total number of tourists staying in location j as the sum of those staying in short-term rentals

and hotels,

$$M_{jt}^T(p_t, a_t) = M_{jt}^{ST}(p_t, a_t) + M_{jt}^H(p_t, a_t). \quad (18)$$

Because the tourist population of a location depends on the vector of prices p_t and amenities matrix a_t , the model endogenizes how the spatial distribution of tourists responds to amenities and prices, but through the STR market. As mentioned above, we consider this reasonable given most of the scope for tourists to switch their accommodation location in response to prices and amenities likely occurs through the more flexible and spatially distributed short-term rental market rather than through the more rigid and spatially clustered hotel sector.

Housing demand from tourists in each location. The impact of tourists on housing demand occurs through STR demand. To express STR demand 17 in units of floor space let f_{jt} denote the average size of a unit in location j . Therefore, STR demand from tourists, in units of floor space is,

$$\mathcal{H}_{jt}^{ST,D}(p_t, a_t) = M_{jt}^{ST}(p_t, a_t) f_{jt}. \quad (19)$$

4.4 Equilibrium

This section defines a stationary equilibrium in which population composition, rents, STR prices, and amenities are endogenously and jointly determined. Before doing so, it is necessary to define a stationary distribution of the population. Consider the type k population law of motion in 13, but written in matrix form,

$$\pi_t^k = \Pi_t^k(r_t, a_t) \pi_{t-1}^k, \quad (20)$$

where each entry in the vector π_t^k is the joint probability of a pair of individual states (j, τ) , while $\Pi_t^k(r_t, a_t)$ is a transition matrix whose entries are the conditional choice probabilities from 12 (and thus depends on rent r_t and amenities a_t).

Definition (Stationary population distribution). Given a vector of rental prices $\mathbf{r} = (r_1, \dots, r_J)$ and a matrix of amenities $\mathbf{a} = [a_1, \dots, a_J]$, a *stationary population distribution over locations and tenure* is a vector $\pi^k(\mathbf{r}, \mathbf{a})$ for each type k that satisfies,

$$\pi^k(\mathbf{r}, \mathbf{a}) = \Pi^k(\mathbf{r}, \mathbf{a}) \pi^k(\mathbf{r}, \mathbf{a}). \quad (21)$$

Notice 21 is simply the stationary version of the law of motion in 20: $\Pi^k(\mathbf{r}, \mathbf{a})$ is the transition matrix implied by rental prices \mathbf{r} and amenities \mathbf{a} , while $\pi^k(\mathbf{r}, \mathbf{a})$ is the stationary population distribution that arises from such a transition matrix. We explicitly denote the population distribution as a function of \mathbf{r} and \mathbf{a} : each entry of the vector $\pi^k(\mathbf{r}, \mathbf{a})$ is the joint probability of an individual state pair (j, τ) , conditional on the aggregate state (\mathbf{r}, \mathbf{a}) . Finally, the stationary distribution implies a stationary type k population count in each location j , which is obtained by summing across

tenure states, i.e., through the stationary version of 14,

$$M_j^k(\mathbf{r}, \mathbf{a}) = \sum_{\tau} \pi^k(\mathbf{r}, \mathbf{a})_{[j, \tau]} M^k \quad \forall k \in \{1, \dots, K\}, \quad (22)$$

where M^k is type- k total population and $\pi^k(\mathbf{r}, \mathbf{a})_{[j, \tau]}$ is entry (j, τ) of vector $\pi^k(\mathbf{r}, \mathbf{a})$.

Definition (Stationary equilibrium). A *stationary equilibrium* is,

1. a vector of long-term rental prices $\mathbf{r} = (r_1, \dots, r_J)$ and a vector of short-term rental prices $\mathbf{p} = (p_1, \dots, p_J)$,
2. a matrix of amenities $\mathbf{a} = [a_1, \dots, a_J]$, where a_j is the vector defined in 2,
3. a stationary population distribution of locals over locations and tenure $\pi^k(\mathbf{r}, \mathbf{a})$ for each type k , which through 22 delivers the type k population count across locations $M^k(\mathbf{r}, \mathbf{a}) = [M_1^k(\mathbf{r}, \mathbf{a}), \dots, M_{J+1}^k(\mathbf{r}, \mathbf{a})]'$ for $k \in \{1, \dots, K\}$,
4. a vector of tourist population $M^{ST}(\mathbf{p}, \mathbf{a}) = [M_1^{ST}(\mathbf{p}, \mathbf{a}), \dots, M_J^{ST}(\mathbf{p}, \mathbf{a})]'$ in short-term rentals, such that,

1. the long-term rental market clears for every location,

$$\underbrace{\frac{\exp(\alpha r_j)}{\exp(\alpha r_j) + \exp(\alpha p_j - \kappa_j)} \mathcal{H}_j}_{\mathcal{H}_j^{LT,S}(r_j, p_j)} = \underbrace{\sum_{k=1}^K M_j^k(\mathbf{r}, \mathbf{a}) f_j^k}_{\mathcal{H}_j^{LT,D}(\mathbf{r}, \mathbf{a})} \quad \forall j,$$

where $\mathcal{H}_j^{LT,S}(\cdot)$ is long-term housing supply defined in 9 and $\mathcal{H}_j^{LT,D}(\cdot)$ is long-term housing demand defined in 15,

2. the short-term rental market clears for every location,

$$\underbrace{\mathcal{H}_j - \mathcal{H}_j^{LT,S}(r_j, p_j)}_{\mathcal{H}_j^{ST,S}(r_j, p_j)} = \underbrace{M_j^{ST}(\mathbf{p}, \mathbf{a}) f_j}_{\mathcal{H}_j^{ST,D}(\mathbf{p}, \mathbf{a})} \quad \forall j.$$

where $\mathcal{H}_j^{ST,S}(\cdot)$ is short-term housing supply, defined residually from the long-term market as in 10, and $\mathcal{H}_j^{ST,D}(\cdot)$ is short-term housing demand from tourists defined in 17,

3. the amenities market clears, by satisfying the mapping defined in 8,

$$a_j = \mathcal{A}(M_j) \quad \forall j,$$

where $M_j \equiv [M_j^1(\mathbf{r}, \mathbf{a}), \dots, M_j^K(\mathbf{r}, \mathbf{a}), M_j^T(\mathbf{p}, \mathbf{a})]'$ is the population composition of location j , consisting of local types 1 through K and tourists. The tourist population includes both those staying in short term rentals as well as the (exogenous) tourist population staying in hotels, defined in 18 as $M_j^T(\mathbf{p}, \mathbf{a}) = M_j^{ST}(\mathbf{p}, \mathbf{a}) + M_j^H(\mathbf{p}, \mathbf{a})$.

A useful interpretation of the equilibrium definition is that conditions 1-2 determine the population distribution of locals and tourists through the clearing of rental markets, for a given distribution of amenities. On the other hand, condition 3 determines the distribution of amenities—as firms enter to clear amenities markets—while taking the population distribution as given. Hence, by combining conditions 1-2 with 3, the population, rents, short-term rental prices, and amenities are jointly and endogenously determined in equilibrium.

In models such as ours, where population composition can have local spillovers, equilibrium uniqueness is not guaranteed. Hence, we use the observed value of prices and amenities as the starting point of our equilibrium solver as a selection rule. In Appendix A.4.2 we show the equilibrium is locally unique under this rule.

5 Estimation

5.1 Defining household heterogeneity

We first classify households into three categories based on modal tenancy status: homeowners, private market renters, and social housing renters. First, ex-ante classification step is motivated by the fact that the average household belongs to its modal category more than 90% of the time, suggesting this margin of adjustment is minor in our context. Second, it allows us to abstract away from the transition between renting and home-ownership. Third, it allows us to separately quantify welfare effects on homeowners and renters in our counterfactual analysis. Hence, we assume tenancy status is determined outside our model and constant over time.

After the first classification step, we classify households further into “types” using a k-means algorithm on demographics. Existing studies typically classify households into groups based on income or skill, while others incorporate additional dimensions, such as race, due to evidence that sorting does not only happen across income levels (Bayer et al., 2016; Davis et al., 2019). When the set of demographic characteristics is large, the practitioner faces a variance-bias trade-off in defining such groups: having more groups captures more heterogeneity but results in fewer observations per group, leading to noisier estimates of choice probabilities. The k-means approach allows us to solve this trade-off in a data-driven manner by exploiting correlations across observables to reduce dimensionality. Further implementation details are in Appendix A.6.1.

Results. Table 2 shows the six household types that result from our classification and summary statistics of their average characteristics. We give each group a label based on how prominent their characteristics are. For example, the “Students” group is characterized by being the youngest and lowest-income, while also being high-skilled and unlikely to have children. Among household types likely to have children, social housing tenants have the lowest income and are less likely to have a Dutch ethnic background. Moving up the income distribution, we have a group of middle-aged homeowners that do not have children, which we label as “Singles”. Next, we have a group of renters that are slightly older and wealthier, but have substantially more children, which we label

Table 2: Summary Statistics by Household Type

Group	Homeowners		Renters		Social Housing Tenants	
	Older Families	Singles	Younger Families	Students	Immigrant Families	Dutch Low Income
Age	44.59	37.84	40.56	28.42	55.12	38.52
Share with Children	0.93	0.12	0.65	0.13	0.53	0.43
Share Low-Skilled	3.20%	2.42%	6.09%	5.40%	99.91%	0.02%
Share Medium-Skilled	3.01%	5.87%	2.28%	11.33%	0.09%	16.95%
Share High-Skilled	93.79%	91.71%	91.65%	83.27%	0.00%	83.02
Share Dutch Indies	6.92%	6.59%	4.12%	4.07%	13.22%	12.41%
Share Dutch	64.41%	58.74%	53.13%	61.44%	24.86%	49.36%
Share Non-Western	18.76%	21.43%	21.64%	19.48%	57.96%	30.37%
Share Western	9.91%	13.23%	21.12%	15.01%	3.96%	7.87%
Household Income (€)	62,031.39	30,611.41	47,441.08	16,821.48	21,243.24	27,714.85
Income Pctl.	77.04	45.49	64.64	23.23	33.41	42.17
Per Capita Income (€)	40,155.65	27,609.21	35,058.39	15,162.83	15,167.45	21,178.13
Income Pctl. per Person	73.42	52.84	65.83	26.34	26.69	42.10
Number of Households	106,388	78,561	105,712	124,112	83,117	174,203

Notes: Table presents the groups resulting from k-means classification on mean demographic characteristics. We report average characteristics across households in each group. “Low”, “medium”, and “high-skilled” correspond to high school or less, vocation/selective secondary education, and college and above, respectively. Group names are chosen to serve as an easy-to-remember label and are not an outcome of the data.

as “Younger Families”. Finally, the highest income group consists of older homeowners likely to have children, which we label as “Older Families”. Overall, the six types vary substantially along income, share with children, and age.

Household types used in estimation and counterfactuals. We estimate the housing demand of local residents, presented in Section 4.3.1, for the first three groups: “Older Families”, “Singles”, and “Younger Families”. The reason for excluding “Students” and the two social housing types is that their residential choices are driven by non-market forces outside the scope of our model. The location choices of “Students” are largely determined by university policy. As for social housing tenants, their units are assigned through a centralized application system.

Despite the exclusion of these three groups in the housing demand estimation, we include all six groups—along with tourists in hotels and in short-term rentals—in the estimation of the amenity supply model described in Section 4.1. The reason is that while residential choices might not be primarily determined by market mechanisms for all groups, as indicated in the prior paragraph, the decisions of firms supplying consumption amenities do take into account all groups regardless of how they make their housing choices. Throughout this section 5, we show our procedure estimates housing demand and amenity supply in separate and independent blocks: estimating

amenity supply only requires neighborhood-level data on population composition, so our sample restriction on the microdata for estimating housing demand does not affect our amenity supply estimates.

Finally, for our counterfactuals we include all six types of locals and tourists as part of our equilibrium definition. Because we do not have preference estimates for students and the two social renter types, we take their location choices as exogenously fixed to levels observed in the baseline data. Given we do not estimate preferences for these groups, we do not make any statements about their welfare effects in our counterfactuals. Our interpretation of keeping the locations of these groups fixed in counterfactuals is that their residential outcomes are determined by an allocation mechanism that does not respond to market forces. Therefore, our counterfactuals should be interpreted as addressing equilibrium responses from the part of the housing market that is determined through market mechanisms.

5.2 Amenities

Re-arranging equation 7 and taking logs, we can rewrite the condition that determines the number of amenities in the sector-location pair sj at time t as,

$$\log N_{sjt} = -\log F_{sjt}(N_{jt}) + \log \left(\sum_k \beta_s^k X_{jt}^k \right), \quad (23)$$

where we define $X_{jt}^k \equiv \phi^k w_t^k M_{jt}^k$ as the total expenditure of the type k population in location j on consumption amenities, and $\beta_s^k \equiv \alpha_s^k / \sigma^s$ as a parameter that dictates how such expenditure is allocated to each amenity sector s . Our microdata allows us to construct X_{jt}^k since income w_t^k is observed in tax returns, population M_{jt}^k is observed in the residential cadaster data, and $1 - \phi^k$, type k 's housing expenditure share, is computed as the ratio of a household's annual expenditure on housing divided by income. Finally, we parameterize the fixed operating cost as follows,

$$F_{sjt}(N_{jt}) = \Lambda_j \Lambda_t R(N_{jt}) \Omega_{sjt},$$

where Λ_j and Λ_t represent location- and year-specific cost shifters, $R(N_{jt})$ is the annual rental price of commercial real estate, and Ω_{sjt} represents any remaining unobservable cost shifters. Because we do not have data on commercial rents, we follow a similar approach as in [Couture et al. \(2021\)](#) and assume that $R(N_{jt}) = N_{jt}^\eta$, where η is the inverse supply elasticity of real-estate. After imposing the fixed cost parameterization, we obtain our estimating equation of amenity supply,

$$\log N_{sjt} = \lambda_j + \lambda_t - \eta \log N_{jt} + \log \left(\sum_k \beta_s^k X_{jt}^k \right) + \omega_{sjt}, \quad (24)$$

where $\lambda_j \equiv -\log \Lambda_j$, $\lambda_t \equiv -\log \Lambda_t$, $\omega_{sjt} \equiv -\log \Omega_{sjt}$. Our main objects of interest are the β_s^k terms, which we infer from the correlation between the sectoral composition of amenities N_{sjt} and the demographic composition of residents M_{jt}^k (which enters 24 through the household-type

composition of amenity expenditure X_{jt}^k).

Identification. First, similar to [Couture et al. \(2021\)](#) we calibrate η . We solve our fully estimated model for a range of $\eta \in [0.39, 1.52]$, which is based on the range of supply elasticities from [Saiz \(2010\)](#). We choose $\eta = 1.52$ since it delivers the best model fit, corresponding to a housing supply elasticity of 0.66, the estimate for San Francisco in [Saiz \(2010\)](#).⁴ In [Appendix A.7.1](#) we show that the main takeaways from our counterfactuals are robust to the full range of $\eta \in [0.39, 1.52]$.

The main identification problem in identifying β_s^k from [24](#) is simultaneity arising from the equilibrium conditions. The expenditures by household type for a given location, X_{jt}^k , are the outcome of residential choices made based on the availability of amenities $N_{s jt}$. Hence, any unobservable firm costs $\omega_{s jt}$ affecting $N_{s jt}$ will also affect M_{jt}^k (and thus X_{jt}^k) in equilibrium. Because $\omega_{s jt}$ is an amenity supply shock, we require instruments that act as amenity demand shifters.

We construct an instrument that shifts population composition, and thus shifts amenity demand differentially across amenity sectors. We use the tax valuation registry to compute the stock of housing units by tenancy status γ in location j , which we define as S_{jt}^γ , where

$$\gamma \in \{\text{owner-occupied, private rental, social housing}\}.$$

We then interact the wages of type k with the housing stock count of their corresponding tenancy status $\gamma(k)$, constructing the following demand shifter,

$$Z_{jt}^k = w_t^k S_{jt}^{\gamma(k)}.$$

The intuition behind our relevance condition is that neighborhoods primarily consisting of social housing units are more likely to be home to households qualifying for social housing assistance, leading to higher expenditure on the specific amenities such households value. The same argument holds for other tenancy types—owner- and renter-occupied units. Our exclusion restriction is therefore,

$$\mathbb{E}[Z_{jt}^k \omega_{s jt} | \lambda_j, \lambda_t] = 0. \tag{25}$$

The above allows for locations with a specific tenancy composition to also have systematically different unobservable fixed costs for the firms supplying amenities. For example, the exclusion restriction would allow for neighborhoods composed mainly of home-owners to have higher commercial real estate rent. However, the exclusion restriction would be violated if the baseline tenancy composition is correlated with *changes* in these unobservable fixed costs. For example, neighborhoods with a higher presence of owner-occupied units could be more likely to tighten local zoning restrictions on services in the future.

⁴We consider San Francisco to be one of the most comparable US cities to Amsterdam in terms of housing supply dynamics: the housing stock of both cities grows at an approximately 1% annual rate (see [San Francisco housing inventory report](#), pg 17, Table 1).

Implementation. We use the six consumption amenities described in section 2: touristic amenities, restaurants, bars, food stores, non-food stores, and nurseries. We simultaneously estimate the parameters in equation 24 for all amenities using GMM. To construct our moments, we interact our instruments with a dummy variable for each amenity s so that $Z_{sjt}^k = \mathbb{1}_s Z_{jt}^k$. We combine Z_{sjt}^k with ω_{sjt} from 24 to construct the term $g^{sk}(\lambda_j, \lambda_t, \beta_s^k)_{sjt} \equiv Z_{sjt}^k \omega_{sjt}$. Hence, the moment conditions that identify the β_s^k coefficients are,

$$\mathbb{E} \left[g^{sk}(\lambda_j, \lambda_t, \beta_s^k)_{sjt} \right] = 0.$$

Fixed effects are identified from the following moment conditions,

$$\mathbb{E} \left[g^j(\lambda_j, \lambda_t, \beta_s^k)_{sjt} \right] = \mathbb{E} \left[\lambda_j \omega_{sjt} \right] = 0, \quad \text{and} \quad \mathbb{E} \left[g^t(\lambda_j, \lambda_t, \beta_s^k)_{sjt} \right] = \mathbb{E} \left[\lambda_t \omega_{sjt} \right] = 0.$$

We stack all moments together to form a final vector of moment conditions:

$$\mathbb{E} \left[g(\lambda_j, \lambda_t, \beta_s^k)_{sjt} \right] = \mathbb{E} [Z_{sjt} \omega_{sjt}] = 0,$$

where $Z_{sjt} \equiv \left[Z_{sjt}^1, Z_{sjt}^1, \dots, Z_{sjt}^K, \lambda_j, \lambda_t \right]'_{s,j,t}$. To ensure our optimization problem is well-defined, we impose the condition $\beta_s^k \geq 0$ for all k, s so that $\log \left(\sum_k \beta_s^k X_{jt}^k \right)$ always exists. Note the β_s^k coefficients are proportional to expenditure shares in our amenity demand model from section 4.1, which naturally have a lower bound at zero. Concretely, we solve for the following constrained optimization problem:

$$\max_{\lambda_j, \lambda_t, \beta_s^k} \hat{g}(\lambda_j, \lambda_t, \beta_s^k)'_{sjt} \hat{W} \hat{g}(\lambda_j, \lambda_t, \beta_s^k)_{sjt} \quad \text{s.t.} \quad \beta_s^k \geq 0 \quad \forall s, k,$$

where $\hat{W} = (Z_{sjt} Z_{sjt}')^{-1}$. Because some estimates lie on the boundary ($\hat{\beta}_s^k = 0$), standard inference does not apply. Hence, we construct standard errors via a Bayesian bootstrap procedure with random weighting (Shao and Tu, 2012).

Results. Our estimates for the β_s^k parameters are shown in Table 3 and broadly align with expected differences in consumption patterns across demographic groups. First, the supply of Nurseries, which is likely the amenity most targeted to locals—and specifically those with children—responds most positively to the three family groups and least to Singles and Tourists. Second, Touristic Amenities respond strongly to Tourists, as expected, but also to Students and Singles that might plausibly have similar consumption patterns to those of Tourists. Third, Restaurants respond mostly to Singles, Students, and Tourists, while Bars respond mostly to Tourists. Fourth, Food Stores estimates are the most homogeneous in that they respond to all groups in similar magnitude. This is consistent with the notion that they provide products (groceries) that are demanded homogeneously across all socioeconomic strata. In terms of magnitudes, our parameter estimates imply an exogenous increase in the number of tourists city-wide by 10% would increase

Table 3: Estimates of amenity supply parameters.

	Touristic Amenities	Restaurants	Bars	Food Stores	Non-Food Stores	Nurseries
Older Families	186.3 [0.0,431.929]	7.374 [0.0,30.401]	0.0 [0.0,0.0]	4.469 [0.0,29.098]	11.359 [0.0,50.577]	980.803*** [368.357,1684.881]
Singles	403.022 [0.0,1816.258]	90.723 [0.0,317.227]	0.0 [0.0,0.0]	102.631 [0.0,327.09]	11.347 [0.0,176.969]	0.0 [0.0,0.0]
Younger Families	0.0 [0.0,0.0]	1.077 [0.0,16.47]	11.365 [0.0,38.983]	52.846** [0.0,123.07]	194.655*** [83.304,331.513]	637.116 [0.0,1326.254]
Students	984.639* [0.0,2014.696]	402.153*** [198.125,673.928]	22.562 [0.0,101.561]	123.078 [0.0,319.632]	2.355 [0.0,1.365]	221.204 [0.0,1759.481]
Immigrant Families	0.122 [0.0,0.0]	5.687 [0.0,49.542]	25.705** [1.126,63.178]	90.549 [0.0,210.37]	127.724** [3.251,331.895]	540.228 [0.0,1740.204]
Dutch Low Income	110.617 [0.0,371.65]	9.908 [0.0,56.277]	0.0 [0.0,0.0]	9.077 [0.0,78.283]	0.0 [0.0,0.0]	0.0 [0.0,0.0]
Tourists	749.072*** [522.649,974.412]	397.274*** [316.498,477.965]	211.571*** [156.406,269.374]	136.337*** [80.554,189.6]	724.223*** [579.264,892.963]	0.0 [0.0,0.0]

Notes: This table reports bootstrap results for coefficients β_s^k from Equation 24 for using a three-way panel of 22 districts in Amsterdam for 2008-2018 over 500 draws. Parameters β_s^k and fixed effects λ_j and λ_t are estimated via GMM, where we restrict parameters to be weakly positive as implied by the microfoundation of the amenity model in Appendix A.3.1. The estimation procedure is outlined in section 5.2 following a Bayesian-bootstrap with random Dirichlet weights. Total expenditure X_{jt}^k is measured in thousands of Euros. Top rows indicate average estimates of the bootstrap samples. Results inside square brackets indicate 95% confidence intervals. We omit estimates of the location and time fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the number of firms in Touristic Amenities, Restaurants, Bars, Food Stores, Non-food Stores, and Nurseries by 2.3%, 0.6%, 2.3%, 0.9%, 2.9%, and 0% respectively.

Observe that of the 42 β_s^k coefficients in Table 3 we have 8 that are zero because our constrained optimization problem places some coefficients on the lower bound of zero. Our interpretation is that if a β_s^k coefficient hits the lower bound, then it means the supply of sector s amenities does not respond to the presence of type k residents. Through the lens of our amenity demand model from section 4.1, this non-response occurs because a coefficient of $\beta_s^k = 0$ implies type k agents do not spend any of their income on sector s amenities. Choosing a lower bound larger than zero would ensure $\beta_s^k > 0$ and thus guarantee positive amenity expenditure shares, but we choose not to do so because this would restrict the parameter space.

Finally, while we cannot directly test the exclusion restriction 25, we can provide suggestive evidence that our instrument is uncorrelated with the unobservable component of fixed costs faced by firms. To the best of our knowledge, the main change in amenity regulations during 2008-2018 was a restriction in the operating hours of restaurant outdoor dining space in residential areas. These restrictions were imposed at the precinct level, a spatial unit larger than the districts at which we implement our estimation. Hence, we can use precinct-year fixed effects to control for the unobservable costs imposed by such regulations on the firms supplying amenities. In Appendix A.7.2 we show that including precinct-year fixed effects does not significantly change our estimates from Table 3. We interpret this as suggestive evidence that our instruments are not significantly correlated with the unobservable fixed costs faced by firms: if they were, the precinct-year fixed

effects would have changed our results significantly, given we know that amenity regulations were indeed modified at the precinct level during this period.

5.3 Housing demand

5.3.1 Housing demand from locals

We estimate preference parameters of local residents using the “Euler Equations in Conditional Choice Probabilities” (ECCP) estimator, building on [Aguirregabiria and Mira \(2010\)](#), [Scott \(2013\)](#), and [Kalouptsi et al. \(2021b\)](#). The method allows us to recover parameters *without* taking a stance on beliefs, computing value functions, or solving the equilibrium, thus reducing computational burden. We proceed to describe the assumptions required for the estimation procedure.

Assumptions. We assume the state variables $\{x, \omega, \epsilon\}$ follow a Markov process, along with the following standard assumptions:

1. **Atomistic agents:** the market-level state ω evolves according to a Markov process that is unaffected by individual-level decisions j or states $\{x, \epsilon\}$,

$$p(\omega'|j, x, \omega, \epsilon) = p(\omega'|\omega).$$

2. **Conditional independence:** the transition density for the Markov process factors as,

$$p(x', \omega', \epsilon'|j, x, \omega, \epsilon) = p_x(x'|j, x, \omega)p_\omega(\omega'|\omega)p_\epsilon(\epsilon').$$

3. **Payoff to outside option:** The flow payoff of living outside the city is normalized to zero, $\bar{u}_t^k(0) = 0 \forall k, t$.⁵

Our ECCP estimator is a two-step estimator. First, we estimate conditional choice probabilities (CCP) directly from the data, using a multinomial logit that exploits information about the conditional state. We show in [Appendix A.6.4](#) that this approach reduces the finite sample bias relative to a non-parametric approach that estimates CCP using frequency estimators. Second, the CCP are plugged into a regression equation that relates differences in the likelihood of two different residential histories to differences in their flow payoffs. To derive this regression equation, we first introduce the concept of renewal actions.


Renewal actions. Two paths of actions are said to exhibit *finite dependence* if after a finite number of periods, the distribution of future states is the same ([Arcidiacono and Miller, 2011](#)). In our model, finite dependence appears whenever two households living in different initial locations, j and j' , choose to move to the same new location \tilde{j} . We call such an action a *renewal action*, because

⁵In a logit model the addition of a constant to all choices leads to the same choice probabilities, which implies that utility levels are not pinned down ([Train, 2009](#)). Hence, we follow common practice in normalizing the payoff of the outside option to zero. Counterfactuals are identified under this normalization if the value of the outside option remains constant ([Kalouptsi et al., 2021a](#)).

the location tenure is reset and the distribution of future states is the same for both households. Because expectations of future payoffs are unobservable to the econometrician, a key difficulty in estimating dynamic models is disentangling variation in current payoffs from continuation values. Renewal actions separate these two components by equalizing continuation values, thus leaving differences in choice probabilities being solely a function of differences in flow payoffs.

Concretely, let $\tau(j, j_{t-1}, \tau_{t-1})$ be the function that maps action j and state $x_t = (j_{t-1}, \tau_{t-1})$ to current location capital. Consider the following path represented by Figure 6: let j and j' denote actions chosen at state $x_t = (j_{t-1}, \tau_{t-1})$, reaching states $x_{t+1} = (j, \tau(j, j_{t-1}, \tau_{t-1}))$ and $x'_{t+1} = (j', \tau(j', j_{t-1}, \tau_{t-1}))$, respectively, and let \tilde{j} be a renewal action chosen at time $t + 1$.

Figure 6: Depiction of path combinations used in the estimation.



From such a path we can derive our main regression equation,

$$Y_{t,j,j',\tilde{j},x_t}^k = u_t^k(j, x_t) - u_t^k(j', x_t) + \beta \left[u_t^k(\tilde{j}, x_{t+1}) - u_t^k(\tilde{j}, x'_{t+1}) \right] + \tilde{v}_{t,j,j',x_t}^k$$

where, $Y_{t,j,j',\tilde{j},x_t}^k \equiv \log \left(\frac{\mathbb{P}_t^k(j, x_t)}{\mathbb{P}_t^k(j', x_t)} \right) + \beta \log \left(\frac{\mathbb{P}_{t+1}^k(\tilde{j}, x_{t+1})}{\mathbb{P}_{t+1}^k(\tilde{j}, x'_{t+1})} \right).$ (26)

On the left hand side, $Y_{t,j,j',\tilde{j},x_t}^k$ is the likelihood of path $\{x_t, x_{t+1}\}$ relative to path $\{x_t, x'_{t+1}\}$. On the right hand side, we have differences in flow payoffs for the two periods in which the paths diverge, and an expectational error we label \tilde{v}_{t,j,j',x_t}^k . We relegate the algebraic derivation of equation 26 to Appendix A.6.3.

The key observation is that at time $t + 1$, when two agents of the same type k choose the renewal action \tilde{j} , they both move to the same individual state and hence their future expected payoffs are the same. Therefore, the value functions from each path cancel each other out at $t + 1$ and disappear from equation 26, which states that differences in the likelihood of path (j_{t-1}, j, \tilde{j}) relative to path (j_{t-1}, j', \tilde{j}) are explained solely by differences in flow utility.

Parametric assumptions on flow utility. We assume the component of flow utility that is common to type k households has the following parametric form,

$$\bar{u}_t^k(j) = \delta_j^k + \delta_t^k + \delta_r^k \log r_{jt} + \delta_a^k \log a_{jt} + \delta_b^k \log b_{jt} + \zeta_{jt}^k, \quad \forall j \neq 0, \quad (27)$$

where preferences over observables such as rent r_{jt} , the vector of consumption amenities a_{jt} , and the vector of exogenous location characteristics b_{jt} vary by type k . We also allow for unobservables

by including fixed effects δ_j^k and δ_t^k , and time-location varying shocks ζ_{jt}^k . To be clear about notation, the coefficients in 27 are all scalars except for δ_a^k and δ_b^k . Recall a_{jt} was defined in equation 2 as a vector that lists the number of firms in each sector s , hence $\delta_a^k \equiv [\delta_1^k, \dots, \delta_s^k, \dots, \delta_S^k]$.

Note location fixed effects δ_j^k capture constant differences in a location j 's payoff with respect to the outside option. Similarly, because δ_t^k only enters the utility of inside locations, it measures how the average attractiveness of those evolves relative to the outside option. After incorporating the individual state variables, the flow payoff for a household i of type k in a location j (inside the city) is,⁶

$$u_i^k(j, x_{it}) = \delta_j^k + \delta_t^k + \delta_r^k \log r_{jt} + \delta_a^k \log a_{jt} + \delta_b^k \log b_{jt} + \zeta_{jt}^k + \delta_\tau^k \log \tau_{it} - MC^k(j, j_{it-1}).$$

The functional form above can be derived as the indirect utility of a household that, conditional on choosing location j , allocates her income optimally across housing and various consumption amenities, as presented in the amenity demand section 4.1 (derivations are in Appendix A.3.2). Importantly, the flow utility parameter for amenity sector s , δ_s^k , maps to the amenity demand parameter α_s^k as follows,

$$\delta_s^k = \left[\alpha_s^k \left(\frac{\phi^k}{\sigma_s - 1} \right) + \gamma_s^k \right] / \sigma_\varepsilon^k, \quad (28)$$

where ϕ^k is the income expenditure share on all consumption amenities, $\sigma_s > 1$ is the substitution elasticity across varieties within amenity sector s , σ_ε^k is the standard deviation of type- k 's idiosyncratic shocks, and γ_s^k accounts for indirect utility spillovers generated by the presence of amenity s beyond utility from direct consumption. Note γ_s^k can be negative if the amenity brings along negative spillovers. For example, if the amenity sector we are considering is bars, the term $\alpha_s^k \left(\frac{\phi^k}{\sigma_s - 1} \right) > 0$ accounts for utility gains from direct consumption at bars, while a negative γ_s^k accounts for the dis-utility from the noise bars bring along. Hence, a negative estimate for δ_s^k can be consistent with a positive valuation for the direct consumption of the amenity ($\alpha_s^k > 0$) if the associated spillovers are sufficiently undesirable (γ_s^k is sufficiently negative). Relatedly, an estimate of zero for β_s^k (which occurs for some sk pairs in Table 3) implies $\alpha_s^k = 0$, but this does not restrict the sign of δ_s^k since γ_s^k can take on any sign.

Implementation. To take 26 to the data we impose the parametric version of flow utility, set $j' = 0$, and impose assumption 3, obtaining our final regression equation,

$$Y_{t,j,\tilde{x}_{it}}^k = \delta_j^k + \delta_t^k + \delta_r^k \log r_{jt} + \delta_a^k \log a_{jt} + \delta_b^k \log b_{jt} + \delta_\tau^k \Delta \tau_{it} - \Delta MC_{it}^k + \zeta_{t,j,x_{it}}^k \quad (29)$$

⁶For the ease of notation we are assuming a deterministic evolution of location capital τ . In Appendix A.6.2, we show how to extend the ECCP equation to stochastic transitions.

where,

$$\begin{aligned}\Delta\tau_{it} &\equiv \tau'(j, x_{it}) - \tau'(0, x_{it}), \\ \Delta MC_{it}^k &\equiv \left[MC^k(j, j_{it-1}) - MC^k(0, j_{it-1}) \right] - \beta \left[MC^k(\tilde{j}, j) - MC^k(\tilde{j}, 0) \right],\end{aligned}$$

and where the last term is the sum of the unobservable time-varying location quality and an expectational error, $\tilde{\zeta}_{t,j,x_{it}}^k = \zeta_{jt}^k + \tilde{v}_{t,j,x_{it}}^k$.

In practice, the locations in our empirical application are Amsterdam's 22 districts ("gebied") and our sample period is 2008 to 2018. We define the outside option as any location outside Amsterdam, and our market as households that have lived in Amsterdam at least once between 2008 and 2020. We set our discount value β equal to 0.85 (De Groot and Verboven, 2019; Diamond et al., 2019). We discretize the location tenure space similar to Rust (1987), defining two bins of location capital: less than three years of tenure or more than four. Appendix A.6.2 shows the technical details of the discretization of the state space. Overall, each group has a total of 46 states per year (23 past locations times two location capital states). We focus on the first three groups—Older Families, Single Households, and Younger Families—because their location choices are primarily driven by market forces, in contrast to households living in social housing or university housing. Note our Older Families and Singles groups are home-owners. In treating their location decisions in the same way as those of renters, we are implicitly assuming they are renting to themselves.

Identification. First, we include the log of the average apartment size and the log of social housing units as additional location characteristics, b_{jt} . We assume the structural error $\tilde{\zeta}_{t,j,x_{it}}^k$ is orthogonal to these characteristics, location fixed effects, tenure and moving costs,

$$\mathbb{E}[\tilde{\zeta}_{t,j,x_{it}}^k | \delta_t^k, \delta_j^k, \log b_{jt}, \Delta\tau_{it}, \Delta MC_{it}^k] = 0 \forall k.$$

Our equilibrium definition implies $\tilde{\zeta}_{t,j,x_{it}}^k$ could include unobservable neighborhood trends that correlate with neighborhood rents r_{jt} and amenities a_{jt} . Therefore, we construct a vector of instruments, Z_{jt} , and estimate demand parameters via two-step optimal GMM with the following moment conditions,

$$\mathbb{E}[Z_{jt} \tilde{\zeta}_{t,j,x_{it}}^k] = 0 \forall k.$$

Recall the error component in equation 29 is the sum of two components: unobservable demand shocks, $\zeta_{t,j,x_{it}}^k$, and expectational errors, $\tilde{v}_{t,j,x_{it}}^k$. Observe that under rational expectations,

$$\mathbb{E}[Z_{jt} \tilde{v}_{t,j,x_{it}}^k | \delta_j, \log b_{jt}, \Delta\tau_{it}, \Delta MC_{it}^k] = 0 \forall k,$$

as $\mathbb{E}[\tilde{v}_{t,j,x_{it}}^k | \mathcal{I}_t] = 0$ for all j, t , and x_{it} , where \mathcal{I}_t is the set of variables realized at time t or before. Therefore, it suffices to find instruments that are orthogonal to unobservable demand shocks,

$$\mathbb{E}[Z_{jt} \zeta_{t,j,x_{it}}^k | \delta_j, \log b_{jt}, \Delta\tau_{it}, \Delta MC_{it}^k] = 0 \forall i, k, j, t.$$

Because we have six amenities, we construct seven instruments in total. Three of those leverage policy changes that can be treated as supply shocks that shift tenancy composition. Concretely, new regulations on the rental market were introduced in 2011, 2015, and 2017 that changed the incentives of landlords to supply their unit as social housing, a private market unit, or as a short-term rental, respectively. See Appendix A.1.1 for full details on each policy change. To introduce spatial variation, we interact a dummy that turns one after the introduction of the policy with the log of the units in the tenancy category exposed to the policy shock in the previous year. Two additional instruments are the log of housing units removed from the housing stock inside location j as well as outside the precinct, which we also interpret as supply shocks.^{7,8} Finally, we follow Bayer et al. (2007) and construct our last two instruments by using variation in changes of social housing units and the average apartment size in other areas of the city outside the precinct. Using these instruments, we find that the first stage regression of a 2SLS estimation has an F-stat of 169.8.

Results. Table 4 shows estimates of the preference parameters in equation 29 over moving costs, location capital, rent, and consumption amenities for our main three groups. All groups exhibit that moving is costly, with costs that increase with distance between past and current location. All households benefit from the accumulation of location capital. Estimates for rent are negative throughout.

Moving on to preferences δ over amenities, note the coefficients $\delta_a^k \equiv [\delta_1^k, \dots, \delta_s^k, \dots, \delta_5^k]$ from equation 29 capture the sum of i) a positive effect from the direct consumption of the amenity, and ii) indirect spillovers that the amenity brings along (e.g., noise from bars) which can be negative. As discussed when we analyzed equation 28, this explains why the signs of the coefficients in Table 4 can be negative.

Moving beyond the interpretation of the sign of the amenity coefficients, comparing the intensity of preferences across household types requires translating the estimates from Table 4 into willingness to pay (WTP) measures. Concretely, the WTP of group k for amenity sector s is computed as the ratio $-\delta_s^k / \delta_r^k$. Using our WTP measure, the coefficients from Table 4 imply that the two family groups are willing to increase their rent by roughly 0.14% in exchange for a 1% increase in the number of nurseries, while the WTP of singles for nurseries is only 0.02%. Restaurants show a positive and significant coefficient for Singles, with a WTP of 0.3% more in rent for a 1% increase in the number of restaurants. For the other groups, the WTP for restaurants is closer to zero. The first two groups perceive a net negative payoff from Touristic Amenities, while the Younger Families exhibit a positive one. In terms of economic magnitude, the first two groups have a WTP of 0.1% and 0.2% more in rent to avoid a 1% increase in Touristic Amenities, respectively. Non-food stores are positively valued by all groups, with the highest WTP for Younger Families. Coefficients

⁷The removal of housing supply can take place in several ways. One way is through demolitions, by government policy or by private initiative. Unfortunately, the microdata do not tell us the agents behind the removal of these units. Another way is that the physical buildings remain in place but lose their status as being habitable for residential purposes, thus effectively removing housing supply.

⁸A precinct (stadsdeel) is a larger spatial unit containing districts. There are seven in Amsterdam.

Table 4: Preference parameter demand estimation results

	Dependent variable: Relative likelihood of renewal paths		
	Older Families	Singles	Younger Families
High Location Capital	0.187*** (0.017)	0.210*** (0.013)	0.264*** (0.014)
Intra-City Moving Cost	-5.916*** (0.015)	-5.337*** (0.011)	-5.384*** (0.012)
Bilateral Moving Cost	-0.067*** (0.000)	-0.059*** (0.000)	-0.041*** (0.000)
In/Out of City Moving Cost	-4.407*** (0.012)	-4.012*** (0.009)	-4.043*** (0.010)
Log Rent	-10.886*** (1.205)	-2.310** (0.999)	-1.964* (1.027)
Log Touristic Amenities	-1.319*** (0.215)	-0.496*** (0.182)	0.317* (0.177)
Log Restaurants	0.288 (0.346)	0.735** (0.305)	-0.280 (0.286)
Log Bars	-0.757*** (0.099)	-0.528*** (0.085)	-0.104 (0.086)
Log Food Stores	-1.695*** (0.327)	-1.216*** (0.281)	-0.540* (0.282)
Log Nonfood Stores	0.427 (0.356)	1.533*** (0.311)	1.383*** (0.302)
Log Nurseries	1.631*** (0.173)	0.044 (0.143)	0.246* (0.147)
<i>N</i>	233772	233772	233772

Notes: This table presents regression results of preference parameters for a dynamic location choice model for 22 districts in Amsterdam for 2008-2019. We estimate preference parameters separately for three groups via two-step optimal GMM. The dependent variable is differences in path likelihoods, after normalizing with respect to the outside option. After this normalization, each type has 46 possible states (23 past locations and two location capital categories), 22 possible actions, and 21 possible renewal actions over 11 years, which leads to 233,772 possible states and two-step path combinations. We omit exogenous controls—the log of social housing units and the log of the average apartment in square meters—for the ease of exposition. Two-step efficient GMM standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

for bars are negative for all groups, suggesting the presence of negative spillovers associated with these amenities, such as noise, that outweigh their consumption benefits. Finally, coefficients for food stores are negative for all groups. Despite the signs being negative, the ordering is fairly intuitive: the WTP of the family groups for food stores is higher (i.e., less negative) than for singles.

5.3.2 Housing demand from tourists

From equation 17 and the normalization of the hotel option's payoff to zero we derive the following regression equation,

$$\log \mathbb{P}_{jt}^{ST} - \log \mathbb{P}_t^H = \delta_j^{ST} + \delta_t^{ST} + \delta_p^{ST} \log p_{jt} + \delta_a^{ST} \log a_{jt} + \zeta_{jt}^{ST}.$$

We use a yearly panel of 95 neighborhoods (wijk) for 2015-2018.⁹ The endogeneity challenge is that prices and amenities are a function of tourists, and therefore correlated with unobservable demand shocks ζ_{jt}^{ST} . In contrast to section 5.3.1, where we deal with this endogeneity problem using an instrumental variable approach, in this part we directly include controls that account for the time-varying quality of locations as perceived by tourists.¹⁰ The reason is that, for this part, we have a direct measure of how tourists perceive the location’s quality through Airbnb review data. We denote $score_{jt}$ as the score that tourists give to the location of the listing they stay in. Our identifying assumption is that there are no unobservables left after controlling for location quality, conditional on the rest of the covariates,

$$\mathbb{E}[\zeta_{jt}^{ST} | \delta_j^{ST}, \delta_t^{ST}, \log p_{jt}, \log a_{jt}, score_{jt}] = 0.$$

Results are shown in Table 5, indicating tourists prefer cheaper locations with more touristic amenities and fewer nurseries. Tourists are willing to pay a 30% higher price for a location with twice as much touristic amenities. We also show estimates of the model without controlling for the score data. When comparing the two specifications we see that coefficients are not statistically different. We interpret this as suggestive evidence that there is little variation coming from time-varying unobservable demand shocks that inform the location choice of tourists that are also correlated with prices and amenities.

Table 5: Tourist demand across locations.

	Dependent Variable: $\log \mathbb{P}_{jt}^{ST} - \log \mathbb{P}_t^H$			
	Baseline		Controlling for reviews	
Log Price Per Guest	-2.725***	(0.820)	-2.660***	(0.759)
Log Touristic Amenities	1.009***	(0.376)	0.838**	(0.394)
Log Restaurants	0.048	(0.259)	0.017	(0.243)
Log Bars	0.051	(0.155)	0.056	(0.164)
Log Food Stores	-0.001	(0.300)	0.037	(0.323)
Log Nonfood Stores	-0.228	(0.417)	-0.185	(0.407)
Log Nurseries	-0.234*	(0.137)	-0.231*	(0.136)
Log Review Scores			4.768	(3.699)
N	371.000		370.000	
R ²	0.529		0.537	

Notes: Table reports estimates of tourists’ preference for neighborhood (wijk-level) characteristics for a static model of location choice, using neighborhood-level data for 2015-2018. Construction of Airbnb supply and prices is described in Appendix A.2. Wijk-level clustered standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

⁹We move to a finer spatial unit in this part of our estimation because the static feature of the tourist choice problem eliminates the issue of poorly defined choice probabilities. We start in 2015 because that is when the Airbnb price data starts (listings, i.e, quantity, data go back before 2015, but prices do not).

¹⁰We prefer this strategy to an IV given how short the panel is and the fact we need to instrument seven endogenous variables, limiting the statistical variation available to identify the parameters.

5.3.3 Connecting amenity demand and supply estimates to equilibrium sorting

Our model predicts several co-location patterns between households types and amenity sectors. If $\delta_s^k > 0$ and $\beta_s^k > 0$, our model predicts positive assortative patterns between type- k households and the amenity s sector. On the contrary, if $\delta_s^k < 0$ and $\beta_s^k = 0$, there is a negative assortative pattern: not only do type- k households move away from locations with amenity s , but their presence does not lead to amenity s entry. These two cases also create incentives of type- k to co-locate together and, thus, acts as an agglomeration force. The intermediate case in which $\delta_s^k < 0$ and $\beta_s^k > 0$ is analytically ambiguous in terms of sorting patterns. Moreover, in such a case the endogeneity of amenities can be thought as a congestion force similar to rent that makes type- k households disperse across locations. Given our range of estimates for β_s^k and δ_s^k in Tables 3 and 4, we should see Older Families positively sort with Nurseries, Nonfood Stores, and Restaurants. Singles should positively sort with Restaurants and Nonfood Stores. Younger Families should positively sort with Nurseries and Nonfood Stores. Finally, Tourists positively sort with Touristic Amenities and Bars.

5.4 Housing supply

Our estimating equation for the supply of long- relative to short-term units is derived by taking the log difference between the two supply choices in equation 9,

$$\log \mathcal{H}_{jt}^{LT,S} - \log \mathcal{H}_{jt}^{ST,S} = \alpha (r_{jt} - p_{jt}) + \kappa_j + \kappa_t + v_{jt},$$

where we have parameterized the operating cost wedge κ_{jt} into location- and time-fixed effects, and v_{jt} stands for any remaining unobservables varying at the jt level. **Instruments.** OLS estimation leads to simultaneity bias from regressing quantities on prices. The solution is an instrument that shifts relative demand for short- versus long-term units. We use predicted tourist demand from a shift-share instrument: the “shift” part of the instrument exploits time variation in worldwide demand for STR as proxied by online search volume (Barron et al., 2021), while the “share” part constructs neighborhood-level exposure to the shift from the historic spatial distribution of touristic attractions. The relevance condition is straightforward: higher predicted demand of tourists raises short- relative to long-term rental prices. The exclusion restriction holds as long as changes in the predicted tourist demand are uncorrelated with changes in the unobservable costs driving landlord’s decisions. Intuitively, the exposure measure is unlikely to be correlated with changes in landlord’s relative costs of renting short- versus long-term.

Table 6: Long-term (LT) relative to short-term (ST) housing supply elasticities

	Dependent variable: $\ln(\text{LT share}) - \ln(\text{ST share})$			
	OLS	IV	IV	IV
LT price-ST price	0.242*	0.287**	0.309**	0.385
	(0.099)	(0.086)	(0.091)	(0.639)
Year FE			X	X
Wijk FE				X
First stage F-stat		65.68	61.62	3.24
Observations	275	275	275	275

Notes: Table reports estimates of landlords' marginal utility of income for a discrete choice model between the short- and long-term rental markets. Data are a panel with 92 locations 2015-2017. Prices are instrumented using a shift-share instrument (Barron et al., 2021) that proxies for demand shocks. Wijk-level clustered standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Results. For this section we end our estimation sample in 2017 because by the end of this year the Amsterdam municipality began to restrict the number of nights that landlords could rent to tourists. We do this to estimate our housing supply elasticity during a period with a stable policy environment, thus avoiding changes in supply that are responding to regulatory changes rather than price changes. Table 6 presents estimates for α , the landlord's marginal utility of income. OLS estimates are downward-biased compared to IV, as expected with simultaneity bias. Our preferred specification is the IV with two-way fixed effects, despite it being less significant than the others, which likely occurs due to little within-neighborhood variation in a short panel. Reassuringly though, IV estimates are fairly stable across all specifications. In terms of economic significance, the results imply that an increase in the gap between STR prices and long-term rental prices of one standard deviation—which is equivalent to a 29% increase—would raise the market share of the short-term relative to the long-term segment by 13.6%.¹¹


5.5 Model fit

To wrap up our estimation section, we show how our model fits the data by simulating a stationary equilibrium for 2017. We assume agents have perfect foresight, we impose the demand shocks $\zeta_j^k = 0$ in steady-state, we take our housing supply estimate from section 5.4, and we calibrate landlords' differential costs to match the STR tourists in each location in 2017. Simulation details are in Appendix A.4.1.

Figures 7-8 plot the simulated endogenous objects—rents and amenities—against the observed objects in the data, showing our model explains a large portion of the variation in rent, STR prices, and amenities by only using variation in observable characteristics, as the unobservable compo-

¹¹This number is computed as $[\exp(\hat{\alpha}) - 1] * 0.29$, with $\hat{\alpha} = 0.385$.

Figure 7: Model fit: Rents and STR prices



Notes: The figure presents scatter plots, linear fit, and 95% confidence intervals of simulated rents and STR prices, against observed rents and prices for 22 districts. Rents are in *Euros/m²* per year. STR prices are average daily prices.

nents of our demand model, ζ_j^k , are set equal to zero. Moreover, the slope of our simulated equilibrium objects and their data counterparts are not statistically different from one. We take these results as evidence that our model, estimated parameters, and equilibrium assumptions are a good approximation of the economic forces reflected in the data.

6 Counterfactuals

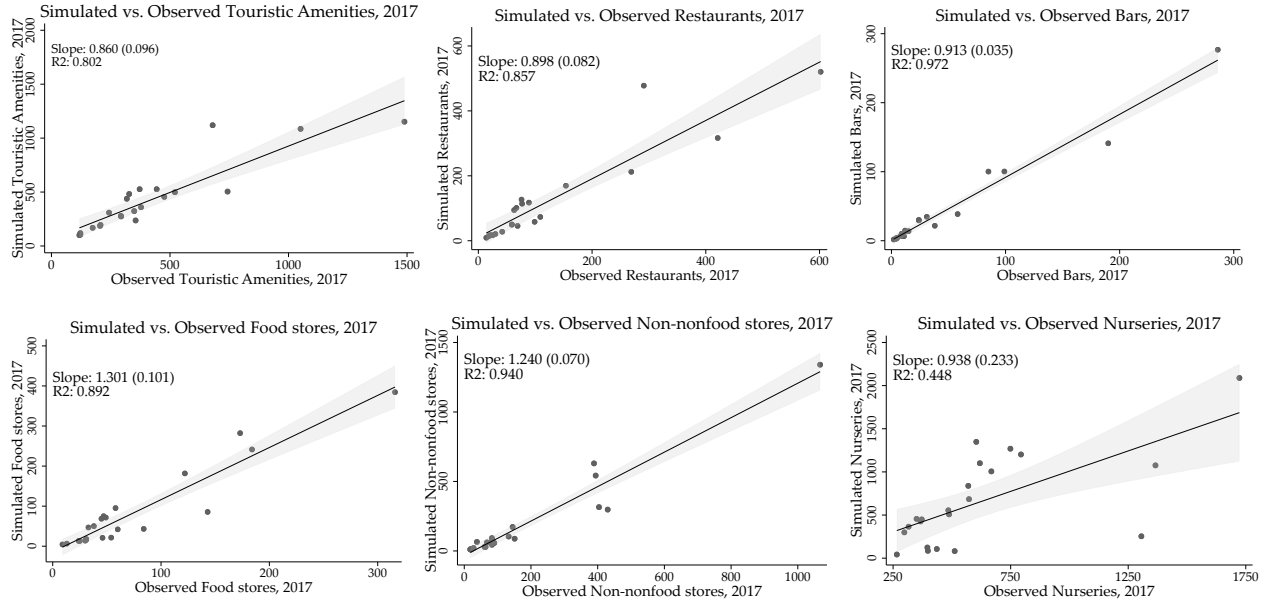
6.1 Role of preference heterogeneity for sorting and inequality

First, we evaluate how preference heterogeneity interacts with the endogeneity of amenities to determine spatial sorting and inequality across residents. We solve the model using the estimates of our baseline heterogeneous preference specification and then compare equilibrium outcomes to those of a homogeneous preference specification. For the homogeneous case we set preference parameters for consumption amenities to the average value across all household types, weighted by the size of groups.¹² We measure sorting with the entropy index, a common measure of residential segregation across household types, with higher values corresponding to more segregation. We measure inequality as the ratio of the highest consumer surplus household (in Euros) to that of the lowest consumer surplus household, with higher values corresponding to more inequality. Our qualitative insights are robust to other measures of inequality.

The left panel of Figure 9 shows segregation is higher when households have heterogeneous preferences for amenities, as they have more neighborhood dimensions along which to sort. The right panel of Figure 9 shows that, despite increased sorting, inequality is lower when preferences are heterogeneous. This empirical result is one of our main takeaways: although heterogeneous

¹²The preferences of tourists are set to the baseline heterogeneous specification in both counterfactuals. In this section, we are only analyzing the role of preference heterogeneity among local residents.

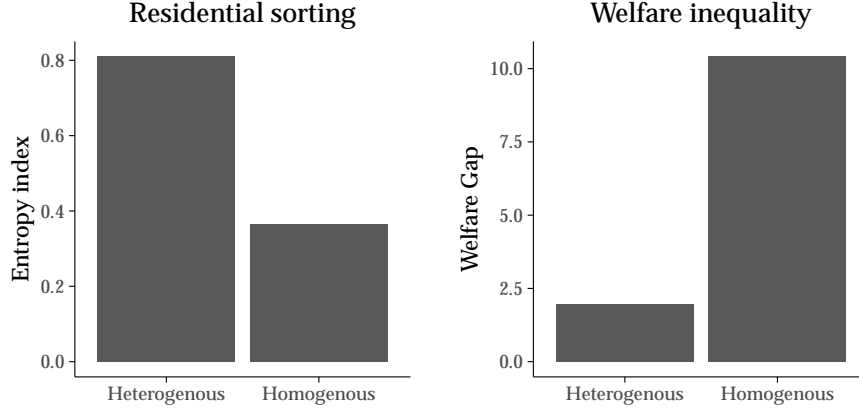
Figure 8: Model fit: Amenities



Notes: The figure presents scatter plots, linear fit, and 95% confidence intervals of the simulated number of amenities against the observed number of amenities for 22 districts. All units are levels.

preferences and endogenous amenities reinforce each other to generate more spatial sorting, they can also reduce welfare inequality across household types. The intuition is that heterogeneous preferences lead to more sorting, which is amplified as amenities respond and make neighborhoods more differentiated. Household inequality can fall if preferences for amenities are heterogeneous because high income groups do not compete with low income groups for the same locations, allowing low income groups to obtain their preferred amenities without having the high income groups bid up their rents. Table 7 conveys the neighborhood differentiation mechanism by showing that all amenities, except one, become more spatially clustered when preferences are heterogeneous, resulting in more differentiated neighborhoods.

Figure 9: Role of preference heterogeneity for spatial sorting and inequality across households.



Notes: The left panel reports the entropy index, a measure of spatial segregation of household types: higher values indicate more segregation (see Appendix A.5.6 for a formal definition). The right panel reports the ratio of the highest consumer surplus household (in Euros) to that of the lowest household: higher values indicate more inequality.

Table 7: Neighborhood differentiation as spatial dispersion of amenities.

Amenity	Gini index for each preference specification		
	Homogenous (HO)	Heterogenous (HE)	HE-HO
Touristic amenities	0.34	0.37	0.03
Restaurants	0.43	0.56	0.13
Bars	0.59	0.66	0.07
Food stores	0.32	0.57	0.25
Non-food stores	0.53	0.67	0.14
Nurseries	0.51	0.43	-0.08

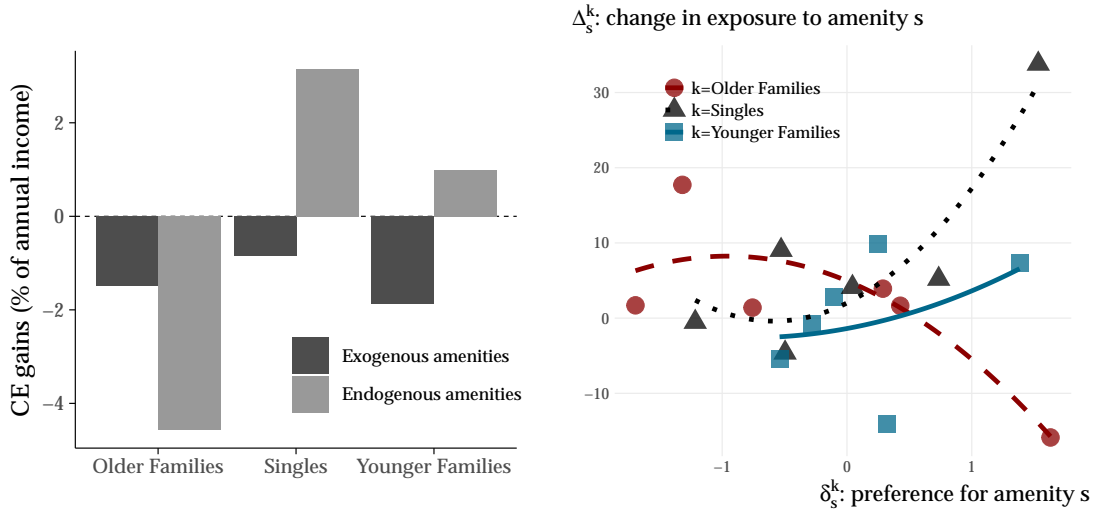
Notes: Columns “Homogeneous” and “Heterogeneous” report the Gini index for each amenity sector: how concentrated the number of establishments in each sector is across locations. Higher values indicate most of the sector’s establishments are clustered in a few locations. Column HE-HO reports the difference between the “Heterogeneous” and “Homogeneous” columns. Positive values in the HE-HO column indicate the spatial distribution of the amenity becomes more clustered across space when preferences are heterogeneous.

6.2 Decomposing welfare effects of the short-term rental industry

In analyzing STR entry, our goal is to disentangle the welfare effects for residents into two components: the increase in rent due to the reduction in housing supply, and changes in amenities due to changes in the composition of amenity demand. To separate these effects, we proceed in three steps. First, we remove the landlords’ STR option and solve for equilibrium rents r_0 and amenities a_0 , which we interpret as the pre-entry equilibrium. Second, we allow landlords to have an STR option but keep amenities fixed at the baseline a_0 , and only solve for rents and STR prices, r and p —the post-entry equilibrium with exogenous amenities. Finally, we allow landlords to have the STR option and simultaneously solve for rents r_1 , STR prices p_1 , and amenities a_1 —the post-entry

equilibrium with endogenous amenities.

Figure 10: Decomposition of welfare effects from STR entry.



Notes: On the left panel, the consumption equivalent (CE) gains on the vertical axes are computed as how much extra income a household must be given in the baseline equilibrium to obtain the same utility as in the counterfactual equilibrium. Therefore, positive values indicate welfare gains due to STR entry. Details in Appendix for A.5.5. On the right panel, the horizontal axis shows preference parameters for amenity sectors. The vertical axis shows the change in exposure to amenity s after STR entry for a type k household, defined as $\Delta_s^k \equiv \sum_j \Delta N_{sj} \times \omega_j^k$, where ΔN_{sj} is the change in sector s amenities in location j after STR entry, weighted by $\omega_j^k = M_j^k / M^k$, location j 's share of the city-wide population of type k before STR entry. Hence, ω_j^k is type k 's exposure to location j .

The left panel of Figure 10 shows the welfare effects from STR entry measured in consumption equivalent (CE) terms: how much extra income a household must be given in the pre-entry equilibrium to be as well off as in the counterfactual post-entry.¹³ Therefore, positive values indicate welfare gains from STR entry. The dark bars show that, under exogenous amenities, every household loses because STR entry reduces housing supply and raises rents. The magnitudes of the losses are similar across household types and equivalent to an income tax between 1-2%.

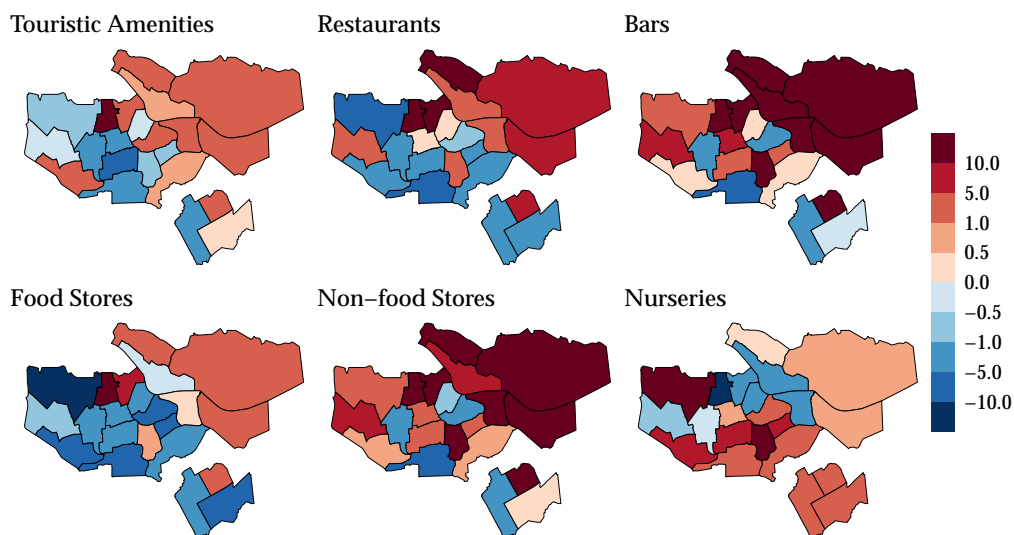
The light bars show the welfare effects when amenities are allowed to endogenously respond to residential composition. The key insight is that while all residents lose due to higher rent, their losses may be compensated or amplified depending on how they value the changes in amenities tourists bring along. Older Families lose more than when amenities were exogenous because on top of facing higher rent they also lose the amenities they value most. On the other hand, Singles and Younger Families now obtain welfare gains because they face an increase in the amenities they like, offsetting losses from higher rent. This mechanism is clearest by looking at the right panel of Figure 10, which plots the correlation between a household type's preferences for amenities and the amenity changes they are exposed to following STR entry. The negative slope for Older

¹³In all cases we take into account differences in home-ownership across household types when computing welfare. Given Table 2, we treat Older Families and Singles as homeowners and Younger Families as renters. Given homeowners rent to themselves, the increase in rent they face due to STR entry is returned to them as landlord income. Details on welfare calculations are in Appendix A.5.

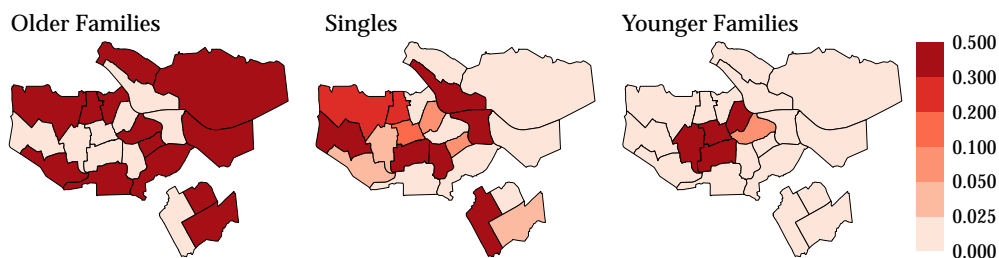
Families implies they are losing access to the amenities they value most. The positive slope for the other groups implies they are gaining access to the amenities they value most.

Figure 11: Effect of STR entry on amenities and baseline distribution of households.

Panel A: Percentage point change in amenities after STR entry.



Panel B: Baseline population shares before STR entry.



Notes: Panel A shows changes in the number of establishments by amenity sector after STR entry. Panel B shows the baseline neighborhood population share of each household type before STR entry, i.e., a measure of exposure to the amenity changes from Panel A. To facilitate comparison between equilibria, we always initialize our equilibrium solver from Appendix A.4.1 with the observed vectors of rents and amenities.

Finally, Figure 11 maps the changes amenities across space and the baseline exposure of each household type to such changes. Note touristic amenities and bars expand the most in locations originally populated by Older Families, and that this group ranks these two amenities among its three least desirable, which explains this group’s negative slope in the right panel of Figure 10.

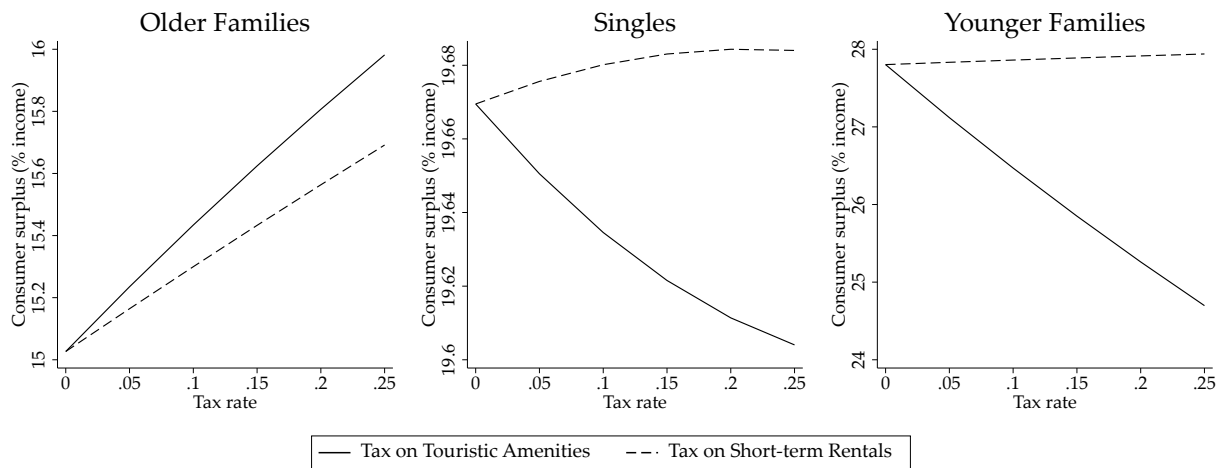
As a final takeaway, note Older Families are the highest-income group and are subject to a welfare loss equivalent to a 5% income tax according to Figure 10. Singles and Younger Families, which are poorer, are subject to welfare gains ranging between 1-3% of their income. In this sense, STR entry is progressive because the higher income group is implicitly taxed at a higher rate. Note

this progressive pattern did not hold with exogenous amenities, since the implicit tax was highest on Younger Families, the middle income group. In this sense, accounting for the endogeneity of amenities can matter for incidence qualitatively, not just quantitatively.

6.3 Policy implications for targeting of amenities

Given our model has both amenity and housing markets, we can compare urban policies that operate separately through each of them. For the purposes of regulating mass tourism, consider two policy levers: a short-term rental (STR) tax or a touristic amenities (TA) tax. The STR tax is a housing policy: its goal is to increase housing supply for locals and improve welfare through rent reductions. The TA tax is an amenity-market policy: it targets certain amenities without directly altering others, but may do so indirectly through equilibrium effects.

Figure 12: Welfare effects: short-term rental tax vs. touristic amenity tax.



Notes: The figure reports consumer surplus (measured as % of income) for each household type under each tax rate.

Figure 12 shows how welfare changes as we gradually increase the tax rate, for each type of tax. First, note that the welfare of all groups is monotonically increasing in the STR tax because the policy reduces rent, and all groups agree they prefer lower rent. However, the rate at which welfare increases is highest for Older Families, since the reduction in STR units also leads to less tourists and touristic amenities, and they especially dislike touristic amenities relative to the other groups.

Second, the shared monotonicity of the tax rate does not hold for the TA tax because the groups disagree on this amenity's desirability. While the welfare of Older Families is increasing, the welfare the other groups is decreasing. This is because Older Families especially dislike touristic amenities and Younger Families value them positively. The case of Singles is more nuanced because they dislike touristic amenities yet somehow lose as the TA tax is increased. The reason is they highly value restaurants, which tend to co-locate with touristic amenities. To see this, note

from our amenity supply estimates in Table 3 that the supply of touristic amenities and restaurants coincide in that they respond strongly to Singles and Tourists. Taxing touristic amenities leads to less Tourists, lowering the supply of restaurants, thus hurting Singles. This highlights the importance of understanding heterogeneity in supply responses of amenities in addition to preference heterogeneity.

To conclude, the incidence of regulating the housing or the amenity market hinges on both preference heterogeneity and supply-side heterogeneity. Therefore, the choice of which policy lever to use depends on the interaction between preference and supply correlations and the distributional objectives of a regulator.

7 Concluding remarks

We study the role of preference heterogeneity over a set of endogenous location amenities in shaping within-city sorting and welfare inequality. To do so, we build a model of residential choice where heterogeneous, forward-looking households consume a bundle of amenities provided by firms in a market for non-tradables. In contrast to work that collapses amenities into a one-dimensional index, we microfound how different consumption amenities arise in equilibrium, endogenizing the extent to which neighborhoods become horizontally differentiated.

Our empirical findings suggest substantial heterogeneity in the preferences of residents for different amenities, as well as in the supply responses of different types of amenities to local demographics. We find that while the endogeneity of amenities reinforces sorting across space, it has ambiguous effects on inequality across households. Concretely, inequality can fall when neighborhoods become horizontally differentiated through the endogenous response of amenities to the sorting of households. Thus, low-income households may sort into neighborhoods only they find desirable, without high-income households bidding up their rents. We also show how the distributional incidence of urban policies depends on heterogeneity on both demand and supply side of the amenities market. While our model is rich in many dimensions, it is tailored to answer a specific set of questions while remaining silent on others. In our concluding remarks, we discuss the limitations of our analysis and potential extensions for future work.

Amenity quality. We do not consider quality differences within an amenity sector because we do not have the firm-level data required to incorporate this dimension. Hence, in our model, amenities are only differentiated horizontally. If we had quality data, then the nature of differentiation across amenities would be a mix of horizontal and vertical dimensions. How would this affect our takeaways? Note our counterfactual in section 6.1 speaks to this because it shows how the degree of horizontal differentiation (measured as the degree of preference heterogeneity) matters for sorting and inequality. To the extent adding amenity quality is a way of dampening horizontal in favor of vertical differentiation, because quality is desired by all groups, we would expect our results to qualitatively change in the direction of the case of homogeneous preferences—there would be less

scope to reduce inequality through sorting and the horizontal differentiation of neighborhoods.

Non-stationarity and transitional dynamics. The role of our model’s dynamic elements is to estimate unbiased preference parameters (Bayer et al., 2016; Traiberman, 2019). To highlight economic mechanisms, our counterfactuals focus on stationary equilibria. The reason is that we are interested in long-run changes that result from the interaction between preference heterogeneity and endogeneity of amenities. It is the cross-sectional correlation between household preferences over amenities and supply responses of amenities to demographics that is at the core of our economic mechanisms. In this sense, it is unclear a-priori that introducing transitional dynamics would significantly change the qualitative nature of our mechanisms, beyond separately quantifying short-versus long-run impacts. Given the stationary analysis already imposes substantial technical complexity, as well as conceptual complexity in understanding how each model ingredient contributes to economic mechanisms, we leave transitional dynamics as an interesting avenue for future research.

Consuming amenities outside the residential location. We assume consumers only access amenities in their residential location. An empirical application that relaxes this assumption requires data on consumption trips across neighborhoods, which we do not have for Amsterdam. Note that our mechanisms are driven by the positive correlation between residential location and amenity consumption. Under our assumption of no commuting to consume, the correlation is perfect. Allowing for commuting would weaken this relationship, but part of the correlation would survive as long as commuting costs depend on distance from home. While we cannot quantify such costs in our setting due to data limitations, smartphone-based evidence from other cities suggests urban residents tend to consume amenities located near their home (Miyauchi et al., 2021; Allen et al., 2021). To the extent this positive correlation between residential location and amenity consumption is also valid in our setting, we expect our main qualitative insights to hold.

References

- Aguirregabiria, V. and Magesan, A. (2013). Euler equations for the estimation of dynamic discrete choice structural models. In: *Structural Econometric Models*, pp. 3–44. Emerald Group Publishing Limited.
- Aguirregabiria, V. and Mira, P. (2010). Dynamic discrete choice structural models: A survey. *Journal of Econometrics* 156(1):38–67.
- Ahlfeldt, G.M., Redding, S.J., Sturm, D.M. and Wolf, N. (2015). The economics of density: Evidence from the Berlin Wall. *Econometrica* 83(6):2127–2189.
- Allen, T., Fuchs, S., Ganapati, S., Graziano, A., Madera, R. and Montoriol-Garriga, J. (2021). Urban Welfare: Tourism in Barcelona.
- Arcidiacono, P. and Miller, R.A. (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* 79(6):1823–1867.
- Barron, K., Kung, E. and Proserpio, D. (2021). The effect of home-sharing on house prices and rents: Evidence from Airbnb. *Marketing Science* 40(1):23–47.
- Bayer, P., Ferreira, F. and McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of Political Economy* 115(4):588–638.

- Bayer, P., McMillan, R., Murphy, A. and Timmins, C. (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica* 84(3):893–942.
- Berry, S., Levinsohn, J. and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica* pp. 841–890.
- Berry, S.T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics* pp. 242–262.
- Botman, M. (2021). Wet toeristische verhuur van woonruimte maakt regulering vakantieverhuur mogelijk.
- Bryan, G. and Morten, M. (2019). The aggregate productivity effects of internal migration: Evidence from Indonesia. *The Journal of Political Economy* 127(5):2229–2268.
- Calder-Wang, S. (2021). The distributional impact of the sharing economy on the housing market. Available at SSRN 3908062 .
- Couture, V., Gaubert, C., Handbury, J. and Hurst, E. (2021). Income growth and the distributional effects of urban spatial sorting. *Review of Economic Studies* .
- Couture, V. and Handbury, J. (2020). Urban revival in America. *Journal of Urban Economics* 119:103267.
- Davis, M., Hartley, D., Gregory, J. et al. (2019). The long-run effects of low-income housing on neighborhood composition. In: *2019 Meeting Papers*, 70. Society for Economic Dynamics.
- Davis, M.A., Gregory, J., Hartley, D.A. and Tan, K.T. (2021). Neighborhood effects and housing vouchers. *Quantitative Economics* 12(4):1307–1346.
- De Groote, O. and Verboven, F. (2019). Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems. *American Economic Review* 109(6):2137–72.
- Diamond, R. (2016). The determinants and welfare implications of US workers’ diverging location choices by skill: 1980–2000. *American Economic Review* 106(3):479–524.
- Diamond, R., McQuade, T. and Qian, F. (2019). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review* 109(9):3365–94.
- Faber, B. and Gaubert, C. (2019). Tourism and economic development: Evidence from Mexico’s coastline. *American Economic Review* 109(6):2245–93.
- Farronato, C. and Fradkin, A. (2018). The welfare effects of peer entry in the accommodation market: The case of airbnb. Tech. rep., National Bureau of Economic Research.
- Fradkin, A., Grewal, E. and Holtz, D. (2018). The determinants of online review informativeness: Evidence from field experiments on Airbnb. Available at SSRN 2939064 .
- García-López, M.À., Jofre-Monseny, J., Martínez-Mazza, R. and Segú, M. (2020). Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona. *Journal of Urban Economics* 119:103278.
- George, L. and Waldfogel, J. (2003). Who affects whom in daily newspaper markets? *Journal of Political Economy* 111(4):765–784.
- Guerrieri, V., Hartley, D. and Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics* 100:45–60.
- Handbury, J. (2021). Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities. *Econometrica* 89(6):2679–2715.
- Hoelzlein, M. (2020). Two-sided sorting and spatial inequality in cities. *Berkeley, mimeograph* .
- Hotz, V.J. and Miller, R.A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies* 60(3):497–529.
- Hsiao, A. (2021). Coordination and commitment in international climate action: Evidence from palm oil. *Unpublished, Department of Economics, MIT* .
- Humlum, A. (2021). Robot adoption and labor market dynamics. *Working Paper* .
- Kalouptsi, M., Scott, P.T. and Souza-Rodrigues, E. (2021a). Identification of counterfactuals in dynamic discrete choice models. *Quantitative Economics* 12(2):351–403.
- Kalouptsi, M., Scott, P.T. and Souza-Rodrigues, E. (2021b). Linear IV regression estimators for structural dynamic discrete choice models. *Journal of Econometrics* 222(1):778–804.
- Koninkrijksrelaties, M.v.B.Z.e. (2015). Wet doorstroming huurmarkt 2015.

- Koninkrijksrelaties, M.v.B.Z.e. (2021). Evaluatie Wet doorstroming huurmarkt 2015.
- Koster, H.R., Van Ommeren, J. and Volkhausen, N. (2021). Short-term rentals and the housing market: Quasi-experimental evidence from Airbnb in Los Angeles. *Journal of Urban Economics* 124:103356.
- McFadden, D. (1974). The measurement of urban travel demand. *Journal of public economics* 3(4):303–328.
- Miyauchi, Y., Nakajima, K. and Redding, S.J. (2021). Consumption access and agglomeration: Evidence from smartphone data. Tech. rep., National Bureau of Economic Research.
- Moretti, E. (2013). Real wage inequality. *American Economic Journal: Applied Economics* 5(1):65–103.
- Mullainathan, S. and Spiess, J. (2017). Machine learning: an applied econometric approach. *Journal of Economic Perspectives* 31(2):87–106.
- Roback, J. (1982). Wages, rents, and the quality of life. *Journal of Political Economy* 90(6):1257–1278.
- Rosen, S. (1979). Wage-based indexes of urban quality of life. *Current issues in urban economics* pp. 74–104.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica: Journal of the Econometric Society* pp. 999–1033.
- Saiz, A. (2010). The geographic determinants of housing supply. *The Quarterly Journal of Economics* 125(3):1253–1296.
- Scott, P.T. (2013). Dynamic discrete choice estimation of agricultural land use. *Toulouse School of Economics Working Paper* 526.
- Shao, J. and Tu, D. (2012). *The jackknife and bootstrap*. Springer Science & Business Media.
- Sheppard, S., Udell, A. et al. (2016). Do Airbnb properties affect house prices? *Williams College Department of Economics Working Papers* 3.
- Su, Y. (2022). The rising value of time and the origin of urban gentrification. *American Economic Journal: Economic Policy* 14(1):402–39.
- Traiberman, S. (2019). Occupations and import competition: Evidence from Denmark. *American Economic Review* 109(12):4260–4301.
- Train, K.E. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press.
- van Perlo, B. (2011). Alle sociale huren 120 euro omhoog. *Het Parool* Section: Voorpagina.
- Vermeulen, W. and Rouwendal, J. (2007). Housing supply and land use regulation in the Netherlands. *Timbergen Institute Discussion Paper*.
- White, M.J. (1986). Segregation and diversity measures in population distribution. *Population index* pp. 198–221.
- Zervas, G., Proserpio, D. and Byers, J.W. (2017). The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry. *Journal of Marketing Research* 54(5):687–705.

A Online Appendix

A.1 Institutional background

A.1.1 Policy changes in the Amsterdam real estate market

Change in Housing Point System (2011): Classification of a unit as social is determined by an annually updated national point system, with units below 143 points being classified as social. Until 2011, the number of points was based solely on the unit’s physical attributes. In response to rental supply scarcity, the Dutch government designated 140 areas nationwide as having a “housing shortage” and implemented a 25-point increase for all rental units in these areas. As a consequence, this policy reduced the supply of social housing units and increased the supply of private rental units. In Amsterdam, it is estimated that 28,000 out of a total of 200,000 social housing units would shift to the private market (van Perlo, 2011).

Decrease in Default Lease Duration (2015): While landlords could historically terminate a rental contract based on certain legal grounds, the contract duration was by default “indefinite.” The only way a landlord could increase rents was with a new lease—to an existing or to a new renter in the private market—or to index the initial lease to inflation. This left little room for landlords to increase rents within a lease. In 2015, a new law, “Wet Doorstroming Huurmarkt 2015”, changed the standard duration of new contracts from indefinite to two years, with options to contract on even shorter duration (Koninkrijksrelaties, 2015). After the initial two years, the landlord had the option to offer the current tenant a new lease with a new price, but which had to be of indefinite duration. As a consequence, landlords had the incentive to find new tenants willing to pay higher prices for an initial two years, rather than renew an existing tenant’s contract indefinitely (Koninkrijksrelaties, 2021), thus increasing private rental market supply.

Regulation of Vacation Rental Properties (2017): Due to the expansion of short-term rentals and tourism, the city of Amsterdam implemented strict regulations on the hospitality sector. First, the policy limited the construction of new hotels (Botman, 2021). Second, the city also required landlords to report all units they rented out as vacation rentals. Third, the city also set a maximum number of nights a property could be rented per year, initially set to sixty nights at the end of 2017 and tightened to only thirty nights in 2019. Together, these laws aimed to lower the incentives to rent short-term to tourists, and thus increase rental supply for locals.

A.2 Data

A.2.1 Residential histories and household characteristics

First, we construct an annual panel of location choices starting in 1995 using the registry (cadaster) data. The cadaster gives us a history of addresses for all individuals in the Netherlands from 1995-2020. For every individual, we pick their modal address each year. In terms of demographics, we keep individuals between 18-70 years old. We also observe country of origin of the household head, which we classify into four broad categories: Dutch, Dutch Indies, Western (OECD), and Non-Western. With regards to skill, we observe the graduation date and degree type for everyone completing a high school degree and beyond for 1999-2020. We classify households according to the highest achieved level of education into low, medium, and high skill for those with high school (VMBO) or less, vocation or selective secondary education (HAVO, VWO, MBO), or college and more (HBO, WO), respectively. Finally, we observe household-level tax returns for 2008-2020 with information on: gross and after-tax income, number of household members, an imputed measure of income per person, and household composition categories. The household composition data allows us to see whether the household has children. For our dynamic location choice estimation sample, we focus on heads of household as identified by the tax data. We keep those households who have lived at least one year in Amsterdam since 1995, household head’s age is between 18-70 years, and have at least one year with reported tax return information.

A.2.2 Housing characteristics, tax appraisal values, and transaction prices

First, for every housing unit we observe the year it was built, the floor area in square meters, a categorical variable about the life stage of the property, and the usage category for 2011-2020. There are 11 usage types: residential, sport, events, incarceration, healthcare, industrial, office, education, retail, and other. There are six types of life-stage categories: constructed, not constructed, in process of construction, in use, demolished, and not in use. We also observe any changes to these characteristics. For example, we can see if a unit previously classified as residential is now considered commercial. With these transitions, we see that virtually no residential units convert to another usage type such as commercial and vice-versa. Given this segmentation, we only keep housing units classified as residential. Moreover, these data inform us about the extent of new construction. On average, new units make 1.2% of the residential housing stock on a yearly basis.

Table 8: Correlation between tax appraisal and transaction values.

	Transaction Value	
WOZ Value	1.069	(0.001)
Constant	-1,798.689	(441.983)
R^2	0.855	
N	128387	

Note: Table shows regression coefficients and fit of transaction values on tax appraisal (WOZ) values at the property level for Amsterdam 2005-2019. Standard errors in parenthesis.

Second, we observe a panel of housing values and characteristics for all properties in the Netherlands from 2006-2019. We observe annual tax appraisal values (WOZ) and geo-coordinates. These data are annually collected by the government to assess every property WOZ value and tax accordingly. The WOZ value of a property is constructed by comparing the value of nearby transacted properties in the neighbourhood and physical housing characteristics like size, house type, and construction year. We can compare WOZ appraisal values to the subset of properties that are transacted to see how well they track market values. Table 8 shows WOZ values correlate almost one-to-one with transaction prices and exhibit a high degree of predictive power. We take this as evidence that WOZ values are informative of market values. These data also contain information about the occupant's tenancy status: homeowner, private renter, or social housing renter. We use these categories to classify households across different segments of the housing market.

A.2.3 Linking households to housing units

We merge the housing unit panel to the household location panel through the property identifier. We can then see tenancy status and number of occupants per unit. We keep housing units with less than six occupants—those below 99th percentile of occupant distribution—to eliminate residential units not inhabited by regular households, such as university student halls or nursing homes.

A.2.4 Rent imputation

Our microdata has information on physical characteristics and tax appraisal values for the universe of housing units in the Netherlands. However, we only observe rents for a subset of units. Because we need an annual panel of housing prices at the neighborhood level, we impute rents using tax valuations.

Table 9: Imputation results.

	In-sample fit				Out-of-sample fit				
	Hedonic Model		Random Forest		Hedonic Model		Random Forest		
	Rental Prices	Price/m ²	Rental Prices	Price/m ²	Rental Prices	Price/m ²	Rental Prices	Price/m ²	
β	1.000 (0.002)	1.005 (0.008)	1.060 (0.002)	1.061 (0.003)	β	1.004 (0.022)	0.957 (0.241)	1.058 (0.007)	1.070 (0.007)
constant	-0.795 (9.364)	-0.091 (0.117)	-73.653 (3.356)	-0.868 (0.037)	constant	5.118 (25.746)	0.764 (0.347)	-73.639 (9.698)	-0.965 (0.107)
R^2	0.636	0.580	0.940	0.940	R^2	0.622	0.554	0.943	0.945
N	11408	11408	11408	11408	N	1268	1268	1268	1268


Notes: Table shows regression coefficients and fit of imputed rental prices on observed rental prices at the property level. We do so for a linear hedonic regression and a random forest, and for two different data samples, the training sample (left panel) to assess in-sample fit, and the testing sample (right panel) to assess out-of-sample fit. Standard errors in parenthesis.

First, we link microdata from the universe of housing units to a national rent survey which contains roughly 13,000 observations of units in the rental market between 2006-2019. We use the matched subset in the rental survey with their tax valuation information to predict rents for housing units that do not appear in the survey but do appear in the property value data as renter-occupied. We keep only properties that are rented in the private rental market and not in social housing. We predict total rental prices and rental prices by square meter on the properties that are classified as private rental units from the tax appraisal data. We use two methods: linear regression and random forest. In both cases we use tax-appraisal values, official categories for measures of quality, total floor area, number of rooms, latitude and longitude coordinates, time fixed effects, and wijk-code fixed effects. We train our algorithms in 90% of the sample and test out-of-sample predictive power in 10% of the sample. For the hedonic linear regression, the in-sample R^2 for total rental prices is 0.637 while the out-of-sample R^2 is 0.629. Similarly, the random forest delivers an in-sample R^2 of 0.813 and out-of-sample R^2 of 0.782. The random forest model has a substantially better performance in terms of predictive power, both in-sample and out-of-sample. Table 9 shows that when regressing imputed on observed rental prices, the random forest also outperforms classic linear regression.

A.2.5 Decreasing hazard rate of moving

Figure 13 shows the hazard rate of moving is decreasing in a household's tenure at the prior residence. This behavior can be rationalized by the inclusion of neighborhood-specific capital that

Figure 13: Probability of changing residence, conditional on past location tenure.



Notes: Figure shows probability of moving out of the current location conditional on the number of years lived in the location. We take averages across individuals who are not social housing residents and across time. Moving probabilities and tenure are constructed using location choice panel derived from the CBS cadaster, described in section A.2.1.

accumulates over time and is lost upon moving.

A.2.6 Housing expenditure shares

With our rent imputation from section A.2.4 we can predict rental prices for all residential units of the city. We compute the share of income spent on housing for households in the private rental market by dividing the predicted rental price by their after-tax income. For households in social housing, we use instead the yearly maximum social housing rent. Finally, we estimate housing expenditures shares by taking the median observation conditional on demographic type and year. These housing expenditure shares map to the term $1 - \phi^k$ in Section A.3.1.

A.2.7 Constructing Airbnb supply and prices

A challenge with the web-scraped Inside Airbnb data is that some listings may be inactive, thus overstating Airbnb supply. To address this we focus on listings that are sufficiently “active”. Using calendar availability data, we say a listing is “active” in month t if it has been reviewed by a guest or its calendar has been updated by its host in month t . Moreover, we want to separately identify listings in which the host lives in the unit and shares it with guests, from those in which there is no sharing. The former does not reduce housing stock for locals, while the latter does. We define a listing as “commercially operated” if it is an entire-home listing, has received new reviews over the past year, and has “sufficient booking activity” such that it is implausible a local is living in

the unit permanently. A listing has “sufficient booking activity” if it satisfies any of the following three conditions:

1. It has been booked over 60 nights in the past year: this is equivalent to over 10 new reviews given an average review rate of 67% (Fradkin, Grewal and Holtz, 2018) and an average stay length of 3.9 nights (source: press.airbnb.com).
2. It shows intent to be booked for many nights over the upcoming year: the listing is available for more than 90 nights over the upcoming year and the “instant book” feature is turned on.
3. It has had frequent updates, reflecting intent to be booked even though it may not have the “instant book” feature turned on: the listing has been actively available for more than 90 nights over the upcoming year and this has happened at least twice within the past year.

A limitation of the data is webscrapes begin in 2015, so we impute listings pre-2015 using calendar and review data. We can only do this for listings that survived up to 2015, therefore our measure of pre- 2015 listings is a conservative lower bound.

A.3 Theory

A.3.1 Derivations of amenity demand

This section derives the amenity demand equation from section 4.1. We model the household decision of how much housing and amenities to consume conditional on living in a specific location. We omit time subscripts unless necessary.

Allocating expenditure between housing and consumption amenities. First, conditional on living in location j , a type k household with Cobb-Douglas preferences chooses how much of its wage w^k to spend on housing H_j and on a bundle of locally available consumption amenities C_j ,

$$\max_{\{H_j, C_j\}} A_j^k H_j^{1-\phi^k} C_j^{\phi^k} \quad \text{s.t.} \quad r_j H_j + P_{C_j} C_j = w^k, \quad (30)$$

where r_j is the rental price, P_{C_j} is the price of the consumption bundle, ϕ^k is the expenditure share parameter for consumption amenities, and A_j^k represents the household’s valuation of the location’s non-market attributes (A_j^k could represent public goods such as noise or pollution). The optimal choice of housing is $H_j^{*,k} = (1 - \phi^k) \frac{w^k}{r_j}$. Therefore, the income left over for amenity consumption is $\phi^k w^k$.

Allocating expenditure across different consumption amenity sectors and varieties. Consumption amenities are classified into sectors indexed $s = 1, \dots, S$ (e.g., “restaurants” is a sector), and firms/varieties are indexed i within each sector (e.g., an Italian restaurant is a firm/variety). The

amenity consumption problem is,

$$\max_{\{q_{isj}^k\}_{is}} C_j^k \text{ s.t. } \sum_{is} p_{isj} q_{isj}^k = \phi^k w^k, \text{ where } C_j^k \equiv \prod_{s=1}^S \left[\left(\sum_{i=1, \dots, N_{sj}} q_{isj}^k \frac{\sigma_s - 1}{\sigma_s} \right)^{\frac{\sigma_s}{\sigma_s - 1}} \right]^{\alpha_s^k}. \quad (31)$$

C_j^k is the amenities bundle, q_{isj}^k is the quantity demanded of variety i in sector-location pair sj , N_{sj} is the number of firms/varieties in the sector-location, and p_{isj} is the price of variety i in sector-location sj . Note C_j^k aggregates consumption amenities across sectors and varieties in a way that is specific to each type k : it implies Cobb-Douglas preferences over amenity sectors (with weights α_s^k , such that $\sum_s \alpha_s^k = 1$) and CES preferences over varieties within an amenity sector (with substitution elasticity $\sigma_s > 1$).¹⁴ Taking first order conditions with respect to q_{isj}^k , and then combining the FOC for two varieties i and i' in the same sector s we obtain,

$$\frac{q_{isj}^k}{q_{i'sj}^k} = \left(\frac{p_{isj}}{p_{i'sj}} \right)^{-\sigma_s}.$$

Furthermore, total expenditure on sector s is $\alpha_s^k \phi^k w^k$ and equal to $\sum_{i \in s} p_{isj} q_{isj}^k$. Using this in the equation above, we obtain type k demand for variety i in sj , i.e., the amenity demand equation from section 4.1 of the main text,

$$q_{isj}^k = \frac{\alpha_s^k \phi^k w^k}{p_{isj}} \left(\frac{p_{isj}}{P_{sj}} \right)^{1-\sigma_s}, \quad \text{with } P_{sj} \equiv \left(\sum_{i=1}^{N_{sj}} p_{isj}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}},$$

A.3.2 Flow utility specification

This section derives the parametric form for flow utility used in estimation in section 5.3.1 of the main text and its connection to the amenity demand parameters.

Indirect utility from housing and amenity demand problem. Given our assumption that marginal costs are constant within a sector-location, the equilibrium of the firm-pricing game is symmetric within a sector-location, thus $p_{isj} = p_{sj} \forall i \in sj$. Hence, consumers buy an equal amount of amenities from every firm within the same sector-location. Type k demand for the individual firm i is,

$$q_{isj}^k = q_{sj}^k = \frac{\alpha_s^k \phi^k w^k}{p_{sj} N_{sj}} \forall i \in sj. \quad (32)$$

To obtain the indirect utility of living in j , we use the equation above to get the optimal amenity

¹⁴Observe that expenditure shares in sector s are identical for households of type j across all neighborhoods j . However, type- k households' utility from consumption amenities differs due to the love-of-variety effect that stems from CES preferences.

bundle C_j^* , which along with the optimal housing choice H_j^* , is substituted in equation 30. To take the indirect utility specification to the data we also reintroduce time subscripts, and impose a flexible form for A_{jt}^k ,

$$A_{jt}^k \underbrace{\frac{w_t^k}{r_{jt}^{1-\phi^k}} \left(\prod_s \left[N_{s jt}^{\frac{1}{\sigma_s-1}} / p_{s jt} \right]^{\alpha_s^k} \right)^{\phi^k}}_{=H_j^{*k} 1-\phi^k C_j^{*\phi^k}} \phi^k, \text{ with } A_{jt}^k \equiv A_j A_t \left(\prod_s N_{s jt}^{\gamma_s^k} \right) b_{jt}^{\alpha_b^k} \tau_t^{\nu^k} \Xi_{jt}^k, \quad (33)$$

and where $\varphi^k \equiv (1 - \phi^k)^{1-\phi^k} (\phi^k)^{\phi^k} \prod_s (\alpha_s^k)^{\alpha_s^k \phi^k}$ is a type- k constant. We assume the valuation of local non-market attributes A_{jt}^k is decomposed as follows: A_j is a fixed location attribute that is unobservable to the econometrician, A_t are unobservable shocks common to all locations in the city, $N_{s jt}^{\gamma_s^k}$ is a utility spillover derived from the nearby presence of amenities beyond the direct consumption itself (which could be dis-utility, such as noise from bars), $\tau_t^{\nu^k}$ is utility from location capital with $\nu^k > 0$, b_{jt} are exogenous time-varying location characteristics that are observable (such as the presence of public housing), and Ξ_{jt}^k are time-varying location attributes that are unobservable. The purpose of 33, especially specifying A_{jt}^k , is to take the theoretical choice problem to the data and be transparent about what the econometrician does and does not observe. Taking logs of 33, and adding a type I EV error ε_{ijt} ,

$$\mu_j^k + \mu_t^k - (1 - \phi^k) \log r_{jt} + \sum_s \left(\frac{\alpha_s^k \phi^k}{\sigma_s - 1} + \gamma_s^k \right) \log N_{s jt} + \alpha_b^k \log b_{jt} + \nu^k \log \tau_t + \zeta_{jt}^k + \varepsilon_{ijt},$$

where $\mu_j^k \equiv \log A_j^k + \log \varphi^k$, $\mu_t^k \equiv \log A_t^k + \log w_t^k$, and $\zeta_{jt}^k \equiv -\phi^k \sum_s \alpha_s^k \log p_{s jt} + \log \Xi_{jt}^k$. Because the level of utility with type I EV errors is not identified, we normalize the variance of the shock to $\frac{\pi^2}{6}$ by dividing the equation above by σ_ε^k ,

$$\delta_j^k + \delta_t^k + \delta_r^k \log r_{jt} + \sum_s \delta_s^k \log N_{s jt} + \delta_b^k \log b_{jt} + \delta_\tau^k \log \tau_t + \zeta_{jt}^k + \varepsilon_{ijt}, \quad (34)$$

where the δ coefficients are the normalized parameters after dividing by σ_ε^k . Finally, to get to the exact flow utility specification from section 5.3.1 of the main text, we define the indirect utility as 34 net of the type I EV shock, we introduce the moving cost, and rewrite $\sum_s \delta_s^k \log N_{s jt}$ in its vector-notation analogue $\delta_a^k \log a_{jt}$,¹⁵

$$u_t^k(j, x_{it}) \equiv \delta_j^k + \delta_t^k + \delta_r^k \log r_{jt} + \delta_a^k \log a_{jt} + \delta_b^k \log b_{jt} + \delta_\tau^k \log \tau_t - MC^k(j, j_{t-1}) + \zeta_{jt}^k.$$

Connection between flow utility parameters and amenity demand parameters. Observe that the flow utility parameters in the last equation above are a function of the parameters of the housing

¹⁵Where $\delta_a^k \equiv [\delta_1^k, \dots, \delta_S^k]$ and $\log a_{jt} \equiv [\log N_{1jt}, \dots, \log N_{Sjt}]'$.

and amenity choice problem,

$$\delta_s^k = \left(\frac{\alpha_s^k \phi^k}{\sigma_s - 1} + \gamma_s^k \right) / \sigma_\varepsilon^k \quad \text{and} \quad \delta_r^k = -(1 - \phi^k) / \sigma_\varepsilon^k.$$

Note the preference parameter for the sector s amenity, δ_s^k , can be positive or negative. The first term, $\frac{\alpha_s^k \phi^k}{\sigma_s - 1}$, is non-negative because the Cobb-Douglas preference parameter for amenity sector s α_s^k is non-negative (consuming the amenity cannot decrease utility). The second term, γ_s^k , can be positive or negative because it measures how the presence of amenity s impacts utility beyond direct consumption through spillovers that can be positive or negative (for example, noise from bars).

A.4 Simulation details

A.4.1 Outline of the equilibrium solver algorithm

We use a nested fixed-point algorithm to solve our model equilibrium. In the inner loop, we solve for the equilibrium vector of long- and short-term rental prices, given a fixed matrix of amenities. In the outer loop, we then solve for equilibrium amenities. The algorithm is as follows: fix parameters $\lambda \in (0, 1)$ and $\delta_r, \delta_p > 0$. The outer loop proceeds as follows for step $t = 1, \dots$

(**O**₁^t) Guess $\mathbf{a}^{(t)}$. The inner loop proceeds as follows for step $g = 1, \dots$

(**I**₁^g) Guess $\mathbf{r}^{(g)}$ and $\mathbf{p}^{(g)}$

(**I**₂^g) Compute excess demand for long- and short-term housing:

$$\mathbf{z}^L(\mathbf{r}^{(g)}, \mathbf{p}^{(g)}, \mathbf{a}^{(t)}) \quad \text{and} \quad \mathbf{z}^S(\mathbf{r}^{(g)}, \mathbf{p}^{(g)}, \mathbf{a}^{(t)})$$

(**I**₃^g) Update prices using excess demands,

$$\begin{aligned} \mathbf{r}^{(g+1)} &= \mathbf{r}^{(g)} + \delta_r \cdot \mathbf{z}^L(\mathbf{r}^{(g)}, \mathbf{p}^{(g)}, \mathbf{a}^{(t)}) \\ \mathbf{p}^{(g+1)} &= \mathbf{p}^{(g)} + \delta_p \cdot \mathbf{z}^S(\mathbf{r}^{(g)}, \mathbf{p}^{(g)}, \mathbf{a}^{(t)}) \end{aligned}$$

(**I**₃^g) Compute $d_{r,p}^{(g)} = \max \left\{ \|\mathbf{r}^{(g+1)} - \mathbf{r}^{(g)}\|_\infty, \|\mathbf{p}^{(g+1)} - \mathbf{p}^{(g)}\|_\infty \right\}$

Iterate until step G such that $d_{r,p}^{(G)} < \epsilon_{r,p}$ for a tolerance level $\epsilon_{r,p} > 0$. Denote,

$$\mathbf{r}^{(et)} \equiv \mathbf{r}^{(G)} \quad \text{and} \quad \mathbf{p}^{(et)} \equiv \mathbf{p}^{(G)}$$

(**O**₂^t) Compute amenities the implied by equilibrium prices from inner loop,

$$\mathbf{a}_{js}^{(et)} = \frac{1}{F_{js} \sigma_s} \left(\sum_{k=1}^K \mathcal{Q}_j^{D,L,k}(\mathbf{r}^{(et)}, \mathbf{a}^{(t)}) \alpha_s^k \alpha_c^k w^k + \mathcal{Q}_j^T(\mathbf{p}^{(et)}, \mathbf{a}^{(t)}) \alpha_s^T \alpha_c^T w^T \right)$$

(O₃^t) Update amenities, $\mathbf{a}^{(t+1)} = (1 - \lambda)\mathbf{a}^{(et)} + \lambda\mathbf{a}^{(t)}$

(O₄^t) Compute $d_a^{(t)} = \|\mathbf{a}^{(t+1)} - \mathbf{a}^{(t)}\|_\infty$

Iterate until step T such that $d_a^{(T)} < \epsilon_a$ for a tolerance level $\epsilon_a > 0$.


Algorithm settings. We construct the amenity supply equation using the estimates from section 5.2. We set the unobservable component of entry costs equal to the residuals of equation 24. For housing demand, we take the estimates from section 5.3, fix the exogenous characteristics of demand at their 2017 level, set unobservable demand shocks ξ_j^k equal to zero (their conditional mean), and sum across groups k to compute aggregate demand for long-term housing. We calibrate the differential costs of short- versus long-term rentals to match the number of STR tourists in each location in 2017. Finally, we start our solver at the observed prices and amenities in 2017. We define convergence when the infinite norm of the excess demand function for the vector of prices and amenities $(\mathbf{r}, \mathbf{p}, \mathbf{a})$ is less than 1E-10.

A.4.2 Local uniqueness of equilibrium

To evaluate the extent of multiplicity, we experiment by perturbing the initial values supplied to the equilibrium solver described in section A.4.1.

Note that given an amenities matrix \mathbf{a} , the equilibrium rent vector \mathbf{r} is unique. Therefore, for our exercise it suffices to vary the initial values of \mathbf{a} . For the perturbation, we first fix the prices to those in the data, $(\mathbf{r}^0, \mathbf{p}^0) = (\mathbf{r}^{Observed}, \mathbf{p}^{Observed})$. Next, we draw an initial amenities matrix \mathbf{a}^0 from a neighborhood around observed amenities, $\mathbf{a}^{Observed}$, as follows: $\mathbf{a}^0 = \mathbf{a}^{Observed} + \mathbf{a}^{Observed} \cdot \epsilon$, where we randomly sample a matrix ϵ from a ring with inner radius ρ and outer radius $\rho + 0.01$, for $\rho = 0, 0.01, \dots, 0.04$. For each ring, we draw 10 different starting values. Figure 14 shows that for any perturbation below $\epsilon = 0.04$ we obtain the same equilibrium. We take this as evidence that at least locally, the equilibrium is unique.

Figure 14: Equilibrium deviations under a range of perturbations.



Notes: The horizontal axes indicate perturbations (ranging from 0% to 4%) of the equilibrium solver's starting point. The vertical axes indicate how the equilibrium that results from the perturbed starting point deviates from the baseline unperturbed equilibrium (in percentage points). Deviations are measured as the mean percentage point gap in equilibrium outcomes across samples (where for each sample we take the median gap in rent, amenities, and STR prices). Values of zero on the vertical axes indicate the perturbation of the starting point leads to the same initial equilibrium. Positive values indicate convergence to a different equilibrium, with higher values indicating further distance from the initial equilibrium.

A.5 Welfare accounting details

A.5.1 Consumer surplus of renters

Following Train (2009), we define consumer surplus as a function of $EV_{j,\tau}^k$ when evaluated at vector (\mathbf{r}, \mathbf{a}) . The expected consumer surplus for a type- k resident is a function of their marginal utility of income v_k and their choice over locations j' :

$$\mathbb{E}[CS_{j,\tau}^k(\mathbf{r}, \mathbf{a})] = \frac{1}{v_k} \mathbb{E}^k \left[\max_{j'} \left(V_{j',j,\tau}^k(\mathbf{r}, \mathbf{a}) + \epsilon_{j'} \right) \right] = \frac{1}{v_k} EV_{j,\tau}^k(\mathbf{r}, \mathbf{a}) + C_k,$$

for some constant C_k . Integrating over the stationary distribution of households over locations, we obtain the following expression for consumer surplus:

$$CS^k(\mathbf{r}, \mathbf{a}) \equiv \frac{1}{v_k} \sum_{j,\tau} EV_{j,\tau}^k(\mathbf{r}, \mathbf{a}) \pi_{j,\tau}^k(\mathbf{r}, \mathbf{a}) + C_k.$$

Following Section A.3.2, the expected value function for group k is,

$$EV_{j,\tau}^k(\mathbf{r}, \mathbf{a}) = \frac{1}{1-\beta} \frac{1}{\sigma^k} \log w^k + f(\mathbf{r}, \mathbf{a}),$$

for some function f and w^k is income of group k . Moreover, we can estimate $\sigma^k = -\frac{1-\phi^k}{\delta_r^k}$ where δ_r^k is the price coefficient and ϕ^k is the housing expenditure share of group k . Hence, the marginal

utility of income for group k can be estimated as:

$$v^k = -\frac{1}{1-\beta} \frac{\delta_r^k}{1-\phi^k} \frac{1}{w^k}. \quad (35)$$

We treat Younger Families as renters, computing their surplus as specified above.

A.5.2 Consumer surplus of home-owners

Some of our household types (Older Families and Singles) are home-owners (i.e., owner-occupiers), whom we assume rent to themselves and receive back rental income. To compute how much rental income, we take a location's average rental income (based on r_j , the long-term rental price per square meter, and $size_j$, the average size of a housing unit) and weight it by the type- k home-owner population,

$$\bar{i}^{L,k}(\mathbf{r}, \mathbf{a}) \equiv \sum_j \frac{Q_j^{D,L,k}(\mathbf{r}, \mathbf{a})}{\sum_j Q_j^{D,L,k}(\mathbf{r}, \mathbf{a})} \cdot r_j \cdot size_j. \quad (36)$$

Consumer surplus of home-owners is the sum of i) their consumer surplus as renters, defined in section A.5.1, and ii) their rental income $\bar{i}^{L,k}(\mathbf{r}, \mathbf{a})$,

$$CS^k(\mathbf{r}, \mathbf{a}) \equiv \frac{1}{v_k} \sum_{j,\tau} EV_{j,\tau}^k(\mathbf{r}, \mathbf{a}) \pi_{j,\tau}^k(\mathbf{r}, \mathbf{a}) + \bar{i}^{L,k}(\mathbf{r}, \mathbf{a}) + C_k.$$

A.5.3 Consumer surplus of tourists

Following Train (2009), the consumer surplus of tourists is given by:

$$\frac{1}{v^T} \log \left(\sum_j \exp(u_j^T(\mathbf{p}, \mathbf{a})) \right) + C_T,$$

where $u_j^T(\mathbf{p}, \mathbf{a}) = \delta_j^S + \delta^S + \delta_p^S \log p_j + \delta_a^S \log a_j + \zeta_j^S$, and $v^T = \sum_j P_j^T(\mathbf{p}, \mathbf{a}) \cdot \frac{\delta_p^S}{p_j}$.

A.5.4 Absentee landlords.

The city-wide average surplus for absentee landlords is,

$$\sum_j \frac{H_j^A}{\sum H_j^A} \left[\frac{1}{\alpha} \log \left(\exp(\alpha p_j + \kappa_j) + \exp(\alpha r_j) \right) + C_L \right],$$

where the term in square brackets is the surplus of the average absentee landlord in location j , weighted by the housing stock owned in each location. We do not include the surplus of absentee landlords in the consumer surplus of residents.

A.5.5 Changes in surplus across counterfactuals

Given two equilibria $(\mathbf{r}_0, \mathbf{a}_0)$ and $(\mathbf{r}_1, \mathbf{a}_1)$, the change in type k consumer surplus is,

$$\Delta \mathbb{E}[CS^k] = \mathbb{E}[CS^k(\mathbf{r}_1, \mathbf{a}_1)] - \mathbb{E}[CS^k(\mathbf{r}_0, \mathbf{a}_0)],$$

where CS^k is defined as in the preceding sections for each household type.

A.5.6 Segregation measure

We use the entropy index (White, 1986) as our measure of segregation. First, we define the entropy index for a single location. Let d_j^k be type k share of location j population—if the type k population in location j is D_j^k , then $d_{jk} \equiv D_j^k / \sum_k D_j^k$. For location j , the entropy index is defined as $v_j \equiv -\sum_{k=1}^K d_j^k \log(d_j^k)$. Next, we define v as the entropy index for the whole city. To do so, we define: $D_j \equiv \sum_{k=1}^K D_j^k$, $D^k \equiv \sum_{j=1}^J D_j^k$, and $D \equiv \sum_{j=1}^J \sum_{k=1}^K D_j^k$, as well as,

$$\hat{v} \equiv -\sum_{k=1}^K \frac{D^k}{D} \log\left(\frac{D^k}{D}\right) \quad \text{and} \quad \bar{v} \equiv \sum_{j=1}^J v_j \frac{D_j}{D} \quad \implies \quad v \equiv \frac{\hat{v} - \bar{v}}{\hat{v}}.$$

Note $v \in [0, 1]$ and higher v means more segregation: v equals 0 if the share of each type in each location is equal to its population share in the whole population, and v equals 1 if each location is occupied by exactly one type.

A.6 Estimation details

A.6.1 Classification by k-means clustering

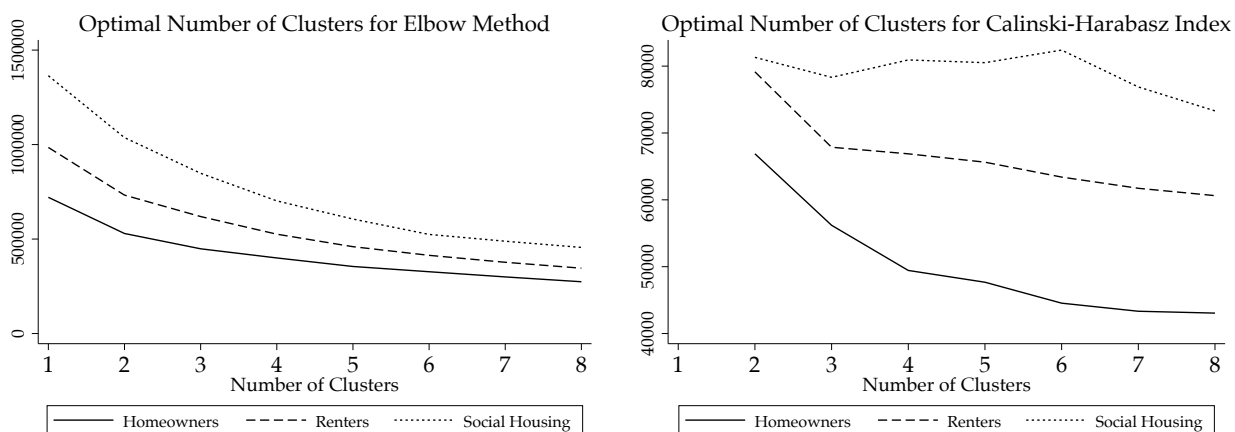
First, given the high persistence in tenancy status, we classify households into three groups based on their modal tenancy status: homeowners, private renters, and renters in social housing. Second, we construct an invariant vector of demographics as follows. For time-varying data—age, disposable income (gross income net of tax), disposable income per person, presence of children—we take averages across years. We standardize all characteristics—skill, region of origin, age, disposable income, disposable income per person, children—because k-means is not invariant to scale and mechanically puts more weight on variables that have larger absolute values. We assign the categorical variables weights of $1/\sqrt{C-1}$, where the number of categories is C , so that each dimension has a weight of 1.¹⁶ We finally run k-means on the transformed vector of demographics.

To choose the number of groups, we use a cross-validation method using two heuristics: the elbow method and the Calinski-Harabasz index. The optimal number of clusters as suggested by the elbow method is pinned down by the largest change of slope in the sum of squared errors

¹⁶That is, for skill, we retain two categories, one that belongs to low skill and one to medium skill. We divide the standardized dummies by $\frac{1}{\sqrt{2}}$. Four country of origin, we set Dutch as the baseline category and divide standardized dummies by $\frac{1}{\sqrt{3}}$.

curve. The Calinski-Harabasz index suggests that the optimal number of clusters is achieved when the ratio of the sum of between-clusters dispersion and of inter-cluster dispersion is maximized. Figure 15 shows the results of these heuristics for the three tenancy groups. For homeowners and private renters both methods suggest an optimal number of two clusters. For social housing renters, the first method suggests two clusters and the second method either two or six clusters. Putting both results together, we choose two as the final number of groups for social housing renters.

Figure 15: Heuristics for k-means classification



A.6.2 Discretization of a continuous state variable

We closely follow Rust (1987). To keep the number of states low, we discretize location tenure in two buckets: $\bar{\tau} = 1$ if $\tau \leq 3$ and $\bar{\tau} = 2$ otherwise. We assume that location tenure evolves using transition probabilities $\mathbb{P}_t(x'_{t+1}|j_t, x_t)$. In practice, we assume $\mathbb{P}_t(\tau_t = 1|j_t, x_t) = 1$ if $j_t \neq j_{t-1}$ and,

$$\mathbb{P}_t(\tau_t = 2|j_t, x_t) = \begin{cases} 1 & , \text{if } j_t = j_{t-1} \text{ and } \tau_{t-1} = 2 \\ p & , \text{if } j_t = j_{t-1} \text{ and } \tau_{t-1} = 1, \end{cases}$$

where p is estimated using a frequency-based estimator.

A.6.3 Constructing the Expected Value Function

The value function is defined as follows:

$$V_t(x, \epsilon) = \max_j \left\{ \mathbb{E}_{x'|j,x} [u_t(x', x)] + \epsilon_j + \beta \mathbb{E}_t [V_{t+1}(x', \epsilon')|j, x, \epsilon] \right\}.$$

Under the assumptions in Section 5.3.1, we define the ex-ante value function as,

$$\mathbb{E}_t[V_{t+1}(x', \epsilon')|j, x, \epsilon] = \int V_{t+1}(x', \epsilon') dF_t(x', \omega_{t+1}, \epsilon'|j, x, \epsilon) \quad (37)$$

$$= \int \left(\int V_{t+1}(x', \epsilon') dF(\epsilon') \right) dF_t(x', \omega_{t+1}|j, x) \quad (38)$$

$$= \int V_{t+1}(x') dF_t(x', \omega_{t+1}|j, x) \equiv EV_t(j, x). \quad (39)$$

We next define the conditional value function:

$$v_t(j, x) = \sum_{x'} \mathbb{P}_t(x'|j, x) (u_t(x', x) + \beta \bar{V}_t(x')) \equiv \bar{u}_t(j, x) + \beta EV_t(j, x).$$

If idiosyncratic shocks are distributed i.i.d. Type I EV, then:

$$\mathbb{P}_t(j|x) = \frac{\exp(v_t(j, x))}{\sum_{j'} \exp(v_t(j', x))}, \quad \text{and} \quad V_t(x) = \log \left(\sum_j \exp v_t(j, x) \right) + \gamma, \quad (40)$$

where γ is Euler's constant. Combining the two previous equations,

$$V_t(x) = v_t(j, x) - \ln(\mathbb{P}_t(j|x)) + \gamma. \quad (41)$$

A key observation is that equation 41 holds for any state x , and any action j .

Toward a demand regression equation. Our demand regression equation's starting point follows Hotz and Miller (1993), by taking differences on equation 40:

$$\ln \left(\frac{\mathbb{P}_t(j|x_t)}{\mathbb{P}_t(j'|x_t)} \right) = v_t(j, x_t) - v_t(j', x_t). \quad (42)$$

Substituting for the choice specific value function,

$$\ln \left(\frac{\mathbb{P}_t(j|x_t)}{\mathbb{P}_t(j'|x_t)} \right) = \bar{u}_t(j, x_t) - \bar{u}_t(j', x_t) + \beta (EV_t(j, x_t) - EV_t(j', x_t)). \quad (43)$$

Following Scott (2013) and Kalouptsi et al. (2021b), the realized expected value $V_t(x')$ can be decomposed between its expectation at time t and its expectational error, where uncertainty is on the aggregate state ω_{t+1} : $V_{t+1}(x') = \bar{V}_t(x') + v_t(x')$. Plugging in everything in equation 43 and

using 41 to replace the continuation values V_{t+1} gives us,

$$\begin{aligned}
\ln \left(\frac{\mathbb{P}_t(j|x_t)}{\mathbb{P}_t(j'|x_t)} \right) &= \sum_x \mathbb{P}(x|j, x_t) u_t(x, x_t) - \sum_{x'} \mathbb{P}(x'|j', x_t) u_t(x', x_t) \\
&\quad + \beta \left[\sum_x \mathbb{P}(x|j, x_t) \bar{V}_t(x) - \sum_{x'} \mathbb{P}(x'|j', x_t) \bar{V}_t(x') \right] \\
&= \sum_x \mathbb{P}(x|j, x_t) u_t(x, x_t) - \sum_{x'} \mathbb{P}(x'|j', x_t) u_t(x', x_t) \\
&\quad + \beta \left[\sum_x \mathbb{P}(x|j, x_t) (V_{t+1}(x) - v_{t+1}(x)) - \sum_{x'} \mathbb{P}(x'|j', x_t) (V_{t+1}(x') - v_{t+1}(x')) \right] \\
&= \sum_x \mathbb{P}(x|j, x_t) u_t(x, x_t) - \sum_{x'} \mathbb{P}(x'|j', x_t) u_t(x', x_t) \\
&\quad - \beta \left[\sum_x \mathbb{P}(x|j, x_t) (v_{t+1}(\tilde{j}, x) - \ln \mathbb{P}_{t+1}(\tilde{j}|x) - v_{t+1}(x)) \right. \\
&\quad \left. - \sum_{x'} \mathbb{P}(x'|j', x_t) (v_{t+1}(\tilde{j}, x') - \ln \mathbb{P}_{t+1}(\tilde{j}|x') - v_{t+1}(x')) \right] \\
&= \sum_x \mathbb{P}(x|j, x_t) u_t(x, x_t) - \sum_{x'} \mathbb{P}(x'|j', x_t) u_t(x', x_t) \\
&\quad - \beta \left[\sum_x \mathbb{P}(x|j, x_t) (v_{t+1}(\tilde{j}, x) - \ln \mathbb{P}_{t+1}(\tilde{j}|x)) \right. \\
&\quad \left. - \sum_{x'} \mathbb{P}(x'|j', x_t) (v_{t+1}(\tilde{j}, x') - \ln \mathbb{P}_{t+1}(\tilde{j}|x')) \right] + \bar{v}_{j,j',x_t},
\end{aligned}$$

where $\bar{v}_{j,j',x_t} \equiv \beta \left(\sum_x \mathbb{P}(x|j, x_t) v_{t+1}(x) - \sum_{x'} \mathbb{P}(x'|j', x_t) v_{t+1}(x') \right)$ is a sum of expectational errors. Observe that if \tilde{j} is a renewal action then:

$$v_{t+1}(\tilde{j}, x) = \bar{u}_{t+1}(\tilde{j}, x) + EV_t(\tilde{j}, 1) = u_{\tilde{j},x,t+1} + \delta_\tau \cdot 1 + MC(\tilde{j}, j) + EV_t(\tilde{j}, 1)$$

for all $x = (j, \tau)$, regardless of τ , where we decompose the per-period utility function, $\bar{u}_{t+1}(\tilde{j}, x)$, into a location specific component, $u_{\tilde{j},x,t+1}$, a location-tenure component δ_τ , and a moving cost component $MC(\tilde{j}, j)$. Substituting and re-arranging,

$$\begin{aligned}
&\ln \left(\frac{\mathbb{P}_t(j|x_t)}{\mathbb{P}_t(j'|x_t)} \right) + \beta \left[\sum_x \mathbb{P}(x|j, x_t) \ln \mathbb{P}_{t+1}(\tilde{j}|x) - \sum_{x'} \mathbb{P}(x'|j', x_t) \ln \mathbb{P}_{t+1}(\tilde{j}|x') \right] \\
&= u_{j,x_t} - u_{j',x_t} + \delta_\tau \left(\sum_x \mathbb{P}(x|j, x_t) \tau(x) - \sum_{x'} \mathbb{P}(x'|j', x_t) \tau(x') \right) \\
&\quad + MC(j, j_{t-1}) - MC(j', j_{t-1}) + \beta \left(MC(\tilde{j}, j) - MC(\tilde{j}, j') \right) + \bar{v}_{j,j',x_t}.
\end{aligned}$$

A.6.4 First-stage estimation of Conditional Choice Probabilities.

We follow a similar procedure as in [Traiberman \(2019\)](#) and [Humlum \(2021\)](#). We depart from their approaches that use a linear probability model and use a multinomial logit on individual decisions to predict choice probabilities for several reasons. First, we can use individual variation. Second, our data reveal that the likelihood of not moving is approximately 85%, while the probability of moving to any other location remains close to zero. This bimodal nature of empirical distribution of choice probabilities suggests that an exponential relationship should be better suited to fit individual decisions compared to a linear model. Third, we find that many predicted probabilities of the linear model lie below zero or above one, a feature that requires an ad-hoc extra censoring step. For every individual i , we observe her individual state at time t $x_{it} = (j_{t-1}, \tau_{t-1})$, where j_{t-1} is the previous location, τ_{t-1} and type $k(i)$, as well as the moving decision variables for all j : $j_{it} = \mathbb{1}\{d(i)_t = j\}$. We define a base outcome 0, and estimate the following multinomial logit model for each group k :

$$\mathbb{P}(j_{it} = j) = \frac{\exp(\lambda_{j,t}^k + \alpha_{j,1}^k \tau_{t-1} + \alpha_{j,2}^k \tau_{t-1}^2)}{1 + \sum_{j'=1}^J \exp(\lambda_{j',t}^k + \alpha_{j',1}^k \tau_{t-1} + \alpha_{j',2}^k \tau_{t-1}^2)}.$$

Monte Carlo simulations. Through a Monte Carlo exercise, we compare the bias in second-stage estimates when first-stage probabilities are predicted with a multinomial logit or with a standard frequency estimator. For our Monte Carlo exercise, we define the period flow utility function as:

$$u_t((d, \tau), x_t) = \alpha \log(r_{dt}) + \sum_s \beta_s \log N_{dst} + \zeta_{dt} + \eta_t + \lambda_d + MC(d, j_{t-1}) + \delta_\tau \tau,$$

with [table 10](#) showing the data generating process of each of the utility components. We also assume that agents have rational expectations. We compute the EV function for each time period as follows. Starting in the last period T , we assume that the economy is in steady-state. We define EV_T as:

$$EV_T(j_T, \tau_T) = \log \left(\sum_d \exp \left(\sum_{x'} \mathbb{P}_T(x'|d, x_T) [u_T(x', x_T) + \beta EV_T(d, x_T)] \right) \right) \quad (44)$$

For $t = 1, \dots, T - 1$, we compute EV_t using backward substitution as follows:

$$EV_t(j_t, \tau_t) = \log \left(\sum_d \exp \left(\sum_{\tau'} \mathbb{P}_{t+1}(x'|d, x_t) [u_{t+1}(d, x_t) + \beta EV_{t+1}(d, x')] \right) \right) \quad (45)$$

Assuming a uniform initial distribution of individuals across states, we simulate each individual forward for 10 time periods. We simulate 10 different samples. We take population sizes of 50 thousand—which roughly corresponds to the size of our groups—and 1 million—which provides insights about convergence properties of large samples. We test two first-stage estimates of conditional choice probabilities: (i) using a multi-nomial logit model and (ii) observed frequencies

where we replace zero shares with a small $\epsilon = 10^{-5}$.

Table 10: Parameters used in simulations

Variables		Parameters	
Name	Distribution (i.i.d.) / Value	Name	Value
u	$N(0, 0.05)$	α	- 0.05
v	$N(0, 0.05)$	β_1, β_2	0.1
ζ	$u + v$	γ_0, γ_1	-0.0025
b_{exo}	$\text{Log}N(0.5, 0.1)$	γ_2	-0.5
r	$0.75 \cdot b_{\text{exo}} + 0.25 \cdot v$	δ	0.1
a_{exo}	$\text{Log}N(1.5, 0.5)$	$N_{\text{households}}$	$\in [5 \cdot 10^5, 10^6]$
a	$0.75 \cdot a_{\text{exo}} + 0.25 \cdot v$	J	24
$\text{dist}(j, j'), j, j' \neq 0; \rho_d$	$\text{Log}N(1, 0.5)$	S	2
λ_j	$N(0, 0.1)$	$\bar{\tau}$	2
λ_t (Perfect foresight)	$N(0, 0.1)$	T	10
λ_t (Rational expectations)	0	tol in EV iteration	10^{-10}

Table 11: Monte Carlo simulations with location fixed effects only and an indicator for high location capital

ζ	Pop (in 10^3)	Prob.	Mean of the absolute value of bias						
			α	β_1	β_2	γ_0	γ_1	γ_2	δ
zero	50	T	1.2E-15	1.9E-16	3.2E-16	1.3E-15	8.2E-19	1.3E-15	1.8E-15
		L	2.3E-02	2.3E-03	4.5E-03	5.7E-02	1.4E-04	3.9E-02	1.0E-01
		F	6.1E-01	1.7E-01	1.7E-01	1.3E+00	2.9E-03	7.8E-01	1.3E+00
	1000	T	8.1E-16	2.5E-16	2.6E-16	1.2E-15	6.5E-19	9.4E-16	1.6E-15
		L	5.6E-03	1.2E-03	8.0E-04	3.3E-02	2.8E-05	3.6E-02	1.0E-01
		F	1.4E-02	3.9E-03	2.6E-03	1.3E-02	5.0E-05	1.4E-02	3.8E-02
exogenous	50	T	2.7E-02	3.8E-03	3.5E-03	7.2E-03	5.0E-06	1.4E-15	2.0E-15
		L	2.7E-02	3.5E-03	5.7E-03	5.9E-02	2.2E-04	2.8E-02	1.0E-01
		F	4.6E-01	2.2E-01	2.7E-01	7.7E-01	2.9E-03	5.2E-01	1.5E+00
	1000	T	4.2E-02	1.5E-02	1.4E-02	7.9E-02	3.6E-04	1.3E-15	2.0E-15
		L	4.4E-02	1.5E-02	1.4E-02	1.0E-01	3.7E-04	3.4E-02	1.0E-01
		F	4.8E-02	1.6E-02	1.4E-02	9.7E-02	4.7E-04	1.3E-02	4.3E-02
endogenous	50	T	2.5E-02	9.9E-03	1.1E-02	9.1E-03	1.3E-05	1.2E-15	2.0E-15
		L	2.8E-02	9.5E-03	1.0E-02	5.6E-02	1.6E-04	4.4E-02	1.0E-01
		F	4.3E-01	2.5E-01	2.4E-01	7.9E-01	2.8E-03	6.5E-01	1.0E+00
	1000	T	3.0E-02	6.8E-03	3.2E-03	1.1E-02	1.0E-05	1.1E-15	1.9E-15
		L	2.8E-02	6.9E-03	3.0E-03	2.5E-02	5.5E-05	3.7E-02	1.0E-01
		F	4.2E-02	8.3E-03	5.5E-03	2.8E-02	1.0E-04	9.2E-03	3.3E-02

Notes: Table presents averaged absolute distance between the estimated parameter and true parameter over 10 random draws of datasets. T represents estimation using the true transition probabilities; L using predicted probabilities by a multinomial logit model; and F using transition probabilities computed based on empirical shares.

The results are presented in Table 11 and reveal that first-stage choice probabilities using a multi-nomial logit model yield a strictly dominant finite sample performance. The gap is most

pronounced in small samples where the likelihood of observing zero flows between states in the data is higher, where the frequency-based estimator uses small but arbitrary values imputed by the researcher, which can be far from the true transition probabilities. The multi-nomial logit approximates the true probabilities well, reducing finite-sample bias in the final estimation stage. As we increase the sample size, the number of observed zero flows diminishes, and we observe convergence of both estimators to the performance of the first-best estimator using the true transition probabilities.

A.7 Robustness exercises

A.7.1 Robustness of real estate supply elasticity

We test how different choices of supply elasticities and their implied congestion parameter η affect our main takeaways. First, Table 12 shows that choosing a supply elasticity equal to San Francisco, as estimated by Saiz (2010), delivers the best model fit in terms of matching the observed distribution of rental prices.

Table 12: Rent fit across a range of supply elasticities

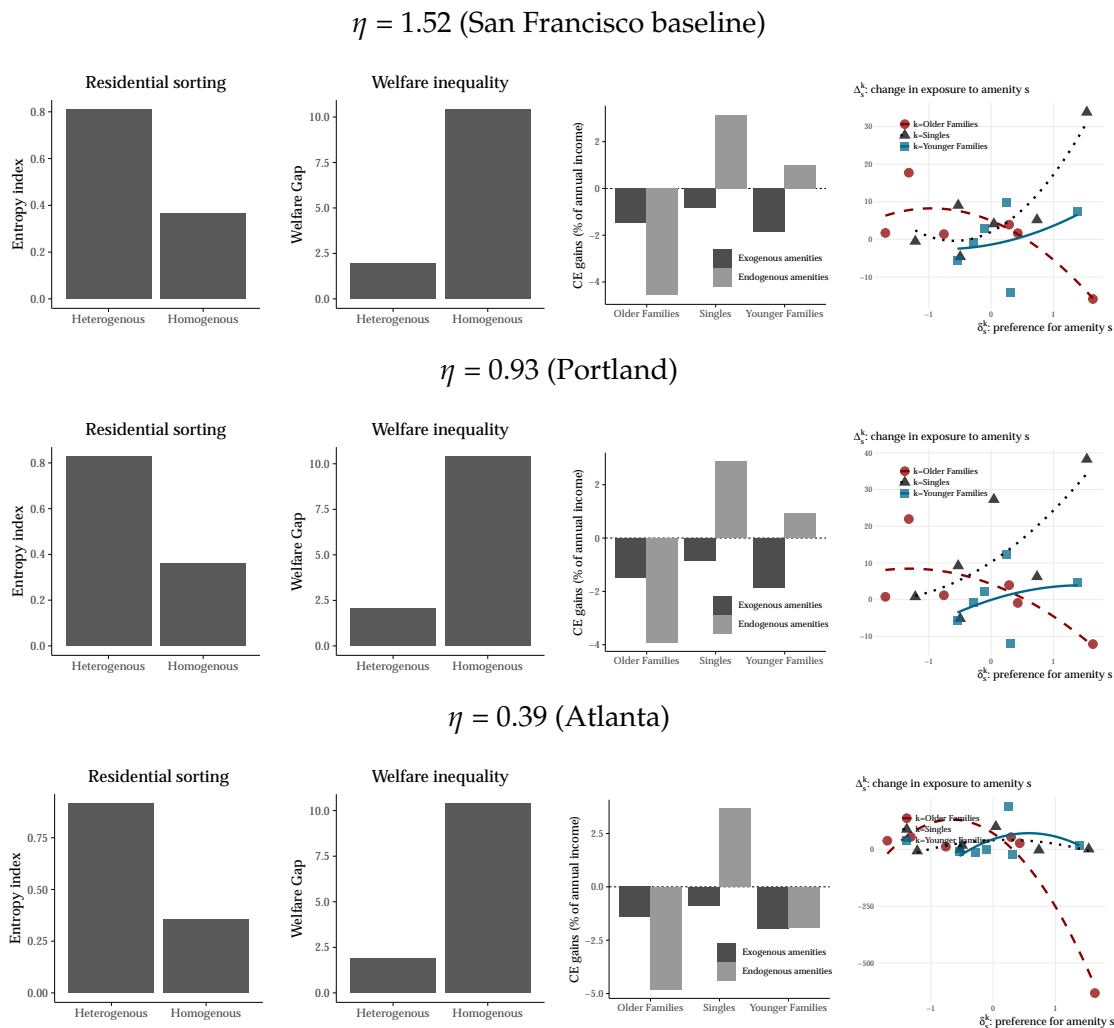
City	Parameters		Rent fit	
	Supply Elasticity	η	R^2	β
San Francisco	0.66	1.52	0.578	1.229
New York	0.75	1.33	0.571	1.229
Boston	0.86	1.16	0.566	1.229
Portland	1.04	0.93	0.557	1.231
Detroit	1.24	0.81	0.549	1.234
Washington DC	1.61	0.62	0.524	1.252
Durham-Raleigh-Chapel Hill	2.11	0.47	0.476	1.291
Atlanta	2.55	0.39	0.473	1.288

Notes: Table presents the R-square and slope of observed rents against our model equilibrium rents. Supply elasticities are from Saiz (2010) and inverted to obtain our amenity congestion parameter η .

Second, Figure 16 shows the key takeaways from our main counterfactuals in sections 6.1-6.2 are robust to different supply elasticities, ranging from our baseline inelastic San Francisco case ($\eta=1.52$) to the highly elastic case of Atlanta ($\eta=0.39$). Figure 16 confirms that for the full range of η , the qualitative insight that preference heterogeneity can lead to more sorting but lower inequality is robust. It also confirms the qualitative insight that all households lose from STR entry due to higher rent, but some are partially compensated by amenity changes depending on how they value the amenities linked to tourism, is robust. In all cases, losses of older families are amplified by endogenous amenities, while those of other groups are compensated. Hence, the choice of η does not make a major difference for the mechanisms in our model, which instead depend on the

correlation between preferences over amenities and amenity supply response across household types.

Figure 16: Robustness of heterogeneity and STR-entry counterfactuals to η .



A.7.2 Robustness of amenity supply estimates to precinct-year fixed effects

We present estimation results for a version of equation 24 from the main text that allows for precinct-year fixed effects:

$$\log N_{sjt} = \lambda_j + \lambda_{p(j)t} - \eta \log N_{jt} + \log \left(\sum_k \beta_s^k X_{jt}^k \right) + \omega_{sjt}, \quad (46)$$

where $p(j)$ indicates the precinct where district j is located. Following the same procedure as in Section 5.2, estimation results are presented in Table 13.

We can test if the difference between the coefficients in Table 13 above and 3 in the main draft, respectively, are statistically indistinguishable. To do so, we can simply check whether confidence

intervals overlap. It is easy to check that for the 95% confidence intervals reported in the tables, we can only reject that one coefficient is statistically different across the two specifications, namely, the coefficient for Tourists on Restaurants.¹⁷ Moreover, to formally test for the difference between the two models, we conduct a joint multiple hypothesis test. The F statistic in that case is given by 0.8197, which is below 1.331—the critical value of an F distribution with 59 and 1319 degrees of freedom at the 5% level. Therefore, we conclude that the two models are not statistically different.

Table 13: Estimates of amenity supply parameters.

	Touristic Amenities	Restaurants	Bars	Food Stores	Non-Food Stores	Nurseries
Older Families	251.65 [0.0,619.269]	15.284 [0.0,46.657]	0.0 [0.0,0.0]	3.042 [0.0,27.438]	18.097 [0.0,85.227]	1268.765*** [497.721,2397.678]
Singles	619.613 [0.0,1989.024]	62.969 [0.0,404.971]	0.0 [0.0,0.0]	160.705 [0.0,530.939]	13.632 [0.0,198.9]	6.78 [0.0,0.001]
Younger Families	0.0 [0.0,0.0]	5.07 [0.0,41.273]	22.95 [0.0,68.748]	96.491*** [12.003,186.982]	331.003*** [161.975,560.076]	1506.579*** [349.155,2863.146]
Students	1891.545** [179.018,3749.272]	713.637*** [354.964,1135.771]	33.078 [0.0,164.968]	174.087 [0.0,561.47]	0.68 [0.0,0.001]	313.761 [0.0,2317.864]
Immigrant Families	0.0 [0.0,0.0]	0.328 [0.0,0.001]	16.918 [0.0,61.799]	64.158 [0.0,170.987]	67.804 [0.0,249.827]	321.911 [0.0,1431.537]
Dutch Low Income	19.156 [0.0,171.662]	1.032 [0.0,16.112]	0.302 [0.0,2.647]	15.877 [0.0,86.375]	0.0 [0.0,0.0]	0.001 [0.0,0.007]
Tourists	1332.384*** [963.408,1732.815]	656.306*** [522.936,817.854]	342.083*** [248.024,443.765]	240.882*** [157.747,308.077]	1151.567*** [880.522,1457.284]	0.0 [0.0,0.0]

Note: Table reports bootstrap results for coefficients β_s^k from Equation 46 for seven population types and six types of services. Parameters β_s^k along with fixed effects λ_j and $\lambda_{p(j)t}$ are estimated via GMM, where we restrict $\beta_s^k \geq 0$. The estimation procedure is outlined in section 5.2 and follows a Bayesian-bootstrap with random Dirichlet weights across 100 draws. Top rows indicate average estimates of the bootstrap samples. Results inside square brackets indicate 95% confidence intervals. We omit estimates of the location and time fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.7.3 Comparison of static and dynamic model estimates

We show how our demand estimates change in the static version of our model: we remove forward-looking behavior (by setting $\beta = 0$) and location capital, i.e., the dynamic state-dependent component of moving costs. We keep the bilateral moving costs since they are a static component of moving costs and are common in static models of migration (Bryan and Morten, 2019).

Table 14 shows the demand estimates in the static model. Table 15 compares the static estimates to our baseline dynamic estimates by performing a t-test of differences for the willingness to pay for amenities—the amenity preference parameters normalized by the rent coefficient. Most of the coefficients are significantly different across specifications, and in several cases even change sign. We take these differences as evidence that failing to account for dynamic considerations can severely bias preference coefficients, in line with the findings of similar studies (Bayer et al., 2016).

¹⁷We reject all statistical differences at the 99% level. We fail to reject four equalities at the 90% level.

Table 14: Preference parameter demand estimation results in the static model.

	Older Families		Singles		Younger Families	
Rent	-17.207***	(4.968)	-11.268***	(4.138)	-14.575***	(4.316)
Tourism Offices	-2.863***	(0.920)	-1.656**	(0.766)	-0.261	(0.799)
Restaurants	2.530	(1.543)	2.990**	(1.285)	1.838	(1.340)
Bars	-0.956**	(0.427)	-0.671*	(0.356)	-0.513	(0.371)
Food Stores	-1.437	(1.359)	0.114	(1.132)	1.120	(1.180)
Nonfood Stores	-1.401	(1.584)	-0.992	(1.320)	-0.225	(1.376)
Nurseries	3.162***	(0.728)	1.629***	(0.607)	2.741***	(0.633)
<i>N</i>	11132		11132		11132	

Notes: Table shows results of preference parameters for a static location choice model for 22 districts for 2008-2019. We estimate preference parameters separately for three groups via GMM. The dependent variable is differences in path likelihoods, after normalizing with respect to the outside option. Each type has 46 possible states, and 22 possible choices over 11 years, leading to 11,132 state-choice combinations. We omit exogenous controls moving costs for ease of exposition. Two-step efficient GMM standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 15: Comparison of dynamic and static estimates.

Group	Amenity	Dynamic		Static		Difference		
		WTP	sd	WTP	sd	Mean	sd	t-test
Older Families	Touristic Amenities	-0.1212	0.0151	-0.1664	0.0396	0.0452	0.0170	2.6667
	Restaurants	0.0265	0.0347	0.1470	0.0896	-0.1206	0.0389	-3.0967
	Bars	-0.0695	0.0091	-0.0556	0.0209	-0.0139	0.0099	-1.4065
	Food Stores	-0.1557	0.0303	-0.0835	0.0746	-0.0721	0.0336	-2.1484
	Nonfood Stores	0.0392	0.0359	-0.0814	0.0980	0.1206	0.0409	2.9518
	Nurseries	0.1498	0.0087	0.1838	0.0267	-0.0340	0.0102	-3.3151
Singles	Touristic Amenities	-0.2148	0.0708	-0.1470	0.0489	-0.0679	0.0699	-0.9702
	Restaurants	0.3183	0.1751	0.2653	0.1313	0.0530	0.1733	0.3057
	Bars	-0.2285	0.0880	-0.0595	0.0268	-0.1690	0.0861	-1.9620
	Food Stores	-0.5266	0.2284	0.0101	0.1017	-0.5368	0.2242	-2.3936
	Nonfood Stores	0.6637	0.3140	-0.0881	0.1256	0.7518	0.3080	2.4412
	Nurseries	0.0190	0.0588	0.1446	0.0276	-0.1256	0.0578	-2.1748
Younger Families	Touristic Amenities	0.1612	0.1676	-0.0179	0.0513	0.1791	0.1641	1.0909
	Restaurants	-0.1426	0.1947	0.1261	0.0907	-0.2687	0.1912	-1.4055
	Bars	-0.0529	0.0379	-0.0352	0.0216	-0.0177	0.0373	-0.4741
	Food Stores	-0.2748	0.1664	0.0769	0.0911	-0.3516	0.1637	-2.1478
	Nonfood Stores	0.7044	0.3846	-0.0154	0.0952	0.7198	0.3763	1.9128
	Nurseries	0.1252	0.0384	0.1881	0.0282	-0.0629	0.0380	-1.6557