Pricing Out the Poor: Income Segregation and Housing Supply Regulation*

Nathan Jones†

November 5, 2025

Click here for the latest version.

Abstract

American cities have been growing more segregated by income. A leading explanation is that the rise in income inequality has increased high-income households' willingness to pay to cluster together. Others have blamed it on ever more restrictive housing supply regulation. I show that these explanations are complementary. To study the interaction of these two forces, I develop a quantitative urban model with a novel margin of endogenous housing supply regulation. Municipalities trade off the profits from new construction with the reduction in value of existing housing that it incurs, and I show that this generates a feedback loop between neighbourhood income and regulation. Municipalities endogenously have stronger incentives to tighten regulation for richer neighbourhoods, pricing out poor households and exacerbating spatial inequality. Quantifying my model with publicly available data, I find that the rise in the college wage premium since 1980 explains 40-86% of the increase in income segregation in New York, Los Angeles, and Chicago, and that 6-29% of this effect comes from the endogenous tightening of regulation.

*I am grateful to my advisors Dávid Nagy and Giacomo Ponzetto for their fantastic support and guidance throughout this project. This paper has also benefited enormously from discussions with Chase Abram, Rodrigo Adão, Milena Almagro, Ayah Bohsali, Thomas Bourany, Simone Cigna, Juanma Castro-Vicenzi, Bruno Conte, Tomás Dominguez-Iino, Elisa Giannone, Esteban Rossi-Hansberg, Matthias Hoelzlein, Olivier Kooi, Vasily Korovkin, Hugo Lhullier, Alberto Martin, Niclas Moneke, Joan Monras, Jordan Rosenthal-Kay, Jaume Ventura, and Shijian Yang. I thank Martin Wiegand for many excellent discussions.

†Department of Economics and Business, Universitat Pompeu Fabra. nathan.jones@upf.edu.

1 Introduction

America's neighbourhoods are staggeringly unequal. The rich live in areas of great wealth, and the poor in areas of concentrated poverty. This pattern of income segregation has grown substantially over the past four decades. Why is this? Some have attributed it to the rise in income inequality, and in particular the growth of the college wage premium (Fogli et al., 2025; Couture et al., 2024). However, others have attributed income segregation to exclusionary housing supply regulation, such as single-family zoning and minimum lot sizes (Trounstine, 2020; Macek, 2024). Such regulation has also increased substantially over the same period (Ganong and Shoag, 2017).

In this paper I show that these narratives, which might seem unrelated, are in fact complementary. First, I incorporate a theory of endogenously formed housing supply regulation into a spatial model of within-city income sorting. My framework generates a feedback loop between neighbourhood income and regulation. As a neighbourhood becomes richer, incentives to block new construction get stronger. Tighter housing supply regulation, in turn, prices out poorer households. My model also provides a way to measure the strength of regulation at the census tract level using publicly available data. Using the model, I then assess the extent to which the rise in the college wage premium since 1980 has increased income segregation precisely by inducing more restrictive regulation in the wealthiest neighbourhoods. I show that this rise explains 40-86% of the increase in income segregation in New York, Los Angeles and Chicago, and that 6-29% of this effect is due to endogenous changes in housing supply regulation.

In my framework, households with different skill types choose whether to live in the city (a Metropolitan Statistical Area) or the rest of the US. Conditional on living in the city, households choose a neighbourhood (a census tract) to live in, and a neighbourhood to work in. Aside from wages and housing prices, they value endogenous neighbourhood amenities that capture positive utility spillovers from living close to high-income households (Diamond, 2016; Almagro and Domínguez-Iino, 2025). Crucially, high-income households are less price-elastic in their demand for housing. That is, housing prices are less important in determining their location choices than they are for poor households (Diamond, 2016; Finlay and Williams, 2022). As a result, where regulation restricts housing supply and pushes up prices, the share of poor households falls (Macek, 2024). This represents the first half of the feedback loop between regulation and neighbourhood income.

The neighbourhoods of the city are divided into municipalities. To build new housing, developers must obtain permits from the relevant municipality. Each permit creates a concentrated benefit and diffuse costs: it allows a developer to make a profit, but this new construction reduces the value of the existing housing stock in the jurisdiction. The municipality assesses each permit application individually, and trades off these two competing effects (Glaeser et al., 2005; Hilber and Robert-Nicoud, 2013). It issues permits in each neighbourhood until it is indifferent between the marginal developer's profit and preventing the decline in inframarginal property values, subject to a neighbourhood-specific weight which I call 'developer power'.

I show that the equilibrium tightness of regulation can be decomposed into two terms: (i) the institutional propensity of municipalities to block new housing, which depends on developer power; and (ii) their economic incentive to do so. This second term represents the decline in property values that new construction induces, and is endogenous to the income composition of the neighbourhood. Since richer households have less elastic demand for housing, new development reduces prices more in high-income neighbourhoods. Municipalities therefore have stronger incentives to restrict housing supply in these neighbourhoods. This completes the feedback loop: regulation makes neighbourhoods richer, and richer neighbourhoods adopt stricter regulation.

A key empirical challenge in solving my model is that municipalities' choices depend on the Jacobian matrix of the housing demand system induced by households' location choices across the entire metropolitan area. That is, they depend on the elasticity of housing demand in one neighbourhood to prices in all neighbourhoods. Unlike in standard quantitative spatial models, it is central to my theory that these elasticities vary with the distribution of income and prices across space. Moreover, the presence of endogenous amenities introduces an 'inner fixed point' into the system of equations defining these elasticities. The large number of neighbourhoods in my counterfactual solutions makes this computation nontrivial. Solving the system that defines the Jacobian by brute force takes several hours. I show, however, that the size of this system can be reduced substantially by exploiting the 'gravity' structure of commuting flows. This renders the problem computationally tractable, and enables me to solve it in under two minutes for any metropolitan area in the US.

I quantify my model with publicly available data. First, I digitise construction cost tables published by a large data aggregator, and map these to data on the mix of housing types in each census tract. This gives me an estimate of marginal construction costs at the neighbourhood level. I combine this with census data on housing prices and households' location choices. Then, using the numerical procedure described above, I invert municipalities' optimality conditions to obtain implied developer power at the census tract level for almost all of the urban United States. This provides a high resolution geography of the institutional propensity for regulation. In simple terms, developer power is inferred to be low in sparsely populated neighbourhoods with high markups of housing prices above marginal construction costs.

Housing supply regulation is notoriously difficult to measure, since it is defined differently in thousands of documents throughout the country. The canonical existing measure, the Wharton Residential Land Use Regulatory Index (Gyourko et al., 2021), is based on a survey of a subset of the municipalities in the US. With minimal data requirements, my measure of developer power instead provides fine spatial detail over a much broader swath of the urban US. It also correlates well with the Wharton index, despite coming from completely different data. Aside from developer power, I recover neighbourhood-specific amenity and productivity terms. Taken together, this set of unobserved fundamentals allows me to conduct counterfactual solutions of the model.

I then use my quantified model to study the rise in income segregation in the US. I show that the rising college wage premium has increased income segregation, and how this has been amplified by endogenous changes in regulation. For ease of exposition, I focus my analysis on the three largest metropolitan areas in the US (New York, Los Angeles, and Chicago), which collectively represent one eighth of the country's population and one fifth of the country's GDP.

Since 1980, the college wage premium has increased by around 75%. I simulate this increase in the model, keeping all other fundamentals fixed, and solve for the resulting equilibria with and without the endogenous adjustment of regulation. This shock causes a shift in average income in all neighbourhoods, and therefore sparks a feedback loop between segregation and regulation. I show that the increased college wage premium explains 40-86% of the observed increase in income segregation, depending on the city, and that the endogenous adjustment of regulation accounts for 6-29% of this effect. At the neighbourhood level, I confirm the key insight of my theory: endogenous regulation makes neighbourhoods that experience income growth become even richer than they otherwise would be.

Related literature: I contribute to two main strands of literature. The first strand studies the determinants of spatial inequality in cities. Among these determinants are geographical heterogeneity (Lee and Lin, 2018; Harari, 2024), urban infrastructure investments (Brinkman and Lin, 2022; Tsivanidis, 2023; Almagro et al., 2024; Bagagli, 2025), or non-homothetic preferences over transport and land consumption (Glaeser et al., 2008). Of particular relevance are the papers that study inequality as a self-reinforcing phenomenon.

Couture and Handbury (2020), Hoelzlein (2023), and Almagro and Domínguez-Iino (2025) focus on the endogenous adjustment of consumption opportunities ('endogenous amenities') as an amplification mechanism for income differences across neighbourhoods. My contribution to this literature is to show that regulation can also act as such an amplification mechanism.

Papers in this literature also study the link between income inequality and income segregation. Fogli et al. (2025) and Couture et al. (2024) both highlight the importance of differential price elasticities of demand for housing among rich and poor households. In each of their models, there is a positive spillover from the rich (concerning human capital accumulation in the case of Fogli et al. (2025), and a neighbourhood-based 'love of variety' effect in Couture et al. (2024)). As income inequality increases, these amenities intensify in certain neighbourhoods, which draws in housing demand. This raises prices, and pushes poor households out of those neighbourhoods. This effect then raises the intensity of the income-based spillovers, and the cycle repeats. Ultimately, segregation is amplified by these feedback loops. I contribute to this work by showing that there is another important mechanism, namely regulation, at play in the link between income inequality and income segregation.

The second broad literature that I contribute to studies housing supply regulation. The last twenty years have seen a great deal of research on this topic, so there are several smaller strands that I contribute to.

One of these strands aims to quantify the effects of housing supply regulation. Glaeser and Gyourko (2002), Ganong and Shoag (2017), and Hsieh and Moretti (2019) discuss how regulation contributes to affordability crises and the misallocation of human capital across cities. There has also been work on the effects of regulation within cities. Rossi-Hansberg (2004) provided early theoretical work in this direction. However, most of the empirical evidence has come more recently, and there is now a robust body of evidence that housing supply regulation lowers density and increases prices at a very local level (Shanks, 2021; Kulka et al., 2022; Hempel, 2023; Nagpal, 2023). Mei (2022) shows how minimum lot size regulations are most hurtful for poorer households, but does not explicitly connect this to spatial sorting. Kulka (2019), Song (2021), and Macek (2024) make this connection, and indeed find that these regulations increase inequality across neighbourhoods. He (2024) and Ma (2025) also find similar results with other types of zoning restrictions.

This literature shows how regulation affects the distribution of economic activity across neighbourhoods. My first contribution is to model this alongside the opposite causal direction; that is, the effect of neighbourhood outcomes on regulation. This allows me to quantify the amplification effects that come from regulation adjusting to sorting changes. My second contribution is methodological. The above papers studying the impact of regulation on sorting across neighbourhoods require confidential parcel-level property records, often from data aggregators such as CoreLogic. My modelling approach is instead able to answer questions about the effect of regulation on sorting with publicly available data.

Another strand of the literature on housing supply regulation focuses on modelling its formation. Fischel (2005), Glaeser et al. (2005), Hilber and Robert-Nicoud (2013), and Ortalo-Magné and Prat (2014) propose theoretical frameworks for understanding the causes of regulation. Two recent papers model endogenous regulation in a general equilibrium setting at the city level. Duranton and Puga (2023) incorporate a model of endogenous permitting costs into a system-of-cities model with human capital spillovers. Their aim is to quantify the aggregate growth effects of loosening restrictions in productive cities. Parkhomenko (2023) models city-wide levels of regulation as arising from a voting process, and studies its effects on the welfare of homeowners and renters. All of these papers emphasise the origin of regulation in the fundamental tension between incumbents (property owners) and newcomers to the area (whose interests are aligned with those of property developers). Metropolitan areas with incumbents that are hurt more by new construction endogenously choose to set higher levels of regulation. At a conceptual level, I continue this approach to modelling housing supply regulation as a conflict between two opposing groups. However, I build on these papers by modelling how the demographics of an area lead to different incentives to regulate, and how this ultimately causes spatial inequality.

This paper also differs from those above by studying the drivers of regulation within metropolitan areas, rather than between them. There is a small empirical literature on this topic. Brooks and Lutz (2019) provide evidence of endogenous regulation leading to persistent effects of infrastructure shocks on built density. I provide a modelling framework to quantify effects like these. Dobbels and Tavakalov (2023) use a minimum lot size reform in Houston to document differential preferences for endogenous amenities, and how this leads to different regulatory decisions at the neighbourhood level. My paper captures this force in a tractable general equilibrium setting, and studies how it ends up exacerbating spatial inequality. Favilukis and Song (2025) document empirical patterns that suggest regulation is driven by municipalities making decisions that benefit their own residents without internalising the externalities that this imposes on the rest of the metropolitan area. At a conceptual level, the conflict between incumbents and outsiders is also at the heart of my modelling framework. I take it further by linking this idea to income sorting and the reinforcement of inequality.

The rest of the paper is organised as follows. In section 2 I describe the model, and highlight a simple version of it that provides intuition for the mechanism I am capturing. Section 3 describes how I bring the model to the data. Section 4 presents the results from counterfactual experiments that decompose observed income sorting increases since 1980.

2 Model

I build a quantitative model of a city in three blocks: i) municipalities that choose housing supply regulation; ii) households with heterogeneous skill types and incomes that choose where to live and where to work; and iii) firms that hire labour and produce a tradable final good.

2.1 Housing Supply Regulation

There is a city with a finite set $\mathcal{N}$ of neighbourhoods, and in each neighbourhood there is a sector of price-taking atomistic housing developers. They produce floorspace h_n, and sell it at price p_n. There is a downward sloping demand curve for floorspace, detailed later in section 2.3, but each developer is too small to internalise this. They produce at marginal cost $c_n(h_n)$. The marginal cost function is neighbourhood specific, reflecting local factors like topography. It is also weakly increasing, reflecting the cost of building at higher density. In low-density neighbourhoods, developers can fill out empty land with single family homes, which is cheap to build. In dense areas, developers must start building upwards rather than outwards, and this entails higher costs.[1]

The city is partitioned into a set $\mathcal{M}$ of municipalities. Municipality m contains the set of neighbourhoods $\mathcal{N}_m \subseteq \mathcal{N}$. Municipalities' only role is to issue permits for housing floorspace. Without loss of generality, I use the same notation, h_n, for floorspace permits and realised floorspace. Municipalities never choose to issue permits that will not be used.

Municipalities assess each marginal permit for a unit of floorspace individually. Each permit affects two parties. The first is the developer who receives the permit, and

[1]In section 3.3 I discuss the particular functional form c_n that I propose and estimate, but this is not central to the theoretical mechanism.

is allowed to produce a marginal unit of floorspace as a result. It makes a profit of $p_n - c_n(h_n)$. The second affected group is a mass of absentee landlords who own the inframarginal floorspace, $h_{n'}$, in each neighbourhood n' in the municipality. When a new unit of floorspace is built in neighbourhood n, the price falls in neighbourhood n' by $\frac{dp_{n'}}{dh_n}$ as housing demand substitutes across neighbourhoods (this is described in section 2.2). The total loss to the owners of existing housing in the neighbourhood is therefore $\sum_{n' \in \mathcal{N}_m} \frac{dp_{n'}}{dh_n} h_{n'}$. The municipality trades off the profits of the marginal developer with the losses of the inframarginal property owners. It issues permits until it is indifferent between these two effects, subject to weights ξ_n:

$$0 = \xi_n \left(p_n - c_n(h_n)\right) + (1 - \xi_n) \left(\sum_{n' \in \mathcal{N}_m} \frac{dp_{n'}}{dh_n} h_{n'} \right) \tag{1}$$

I refer to the weights ξ_n as 'developer power', since they represent the influence that developers in neighbourhood n have relative to property owners in the municipality's decision making. As $\xi_n \to 1$, one can see that the municipality will build housing until the price is equal to the marginal cost of building in that neighbourhood. If, instead, $\xi_n \to 0$, the municipality will never want to permit any new housing. This formulation captures the fundamental conflict at the heart of housing supply regulation. Fischel (2005) highlights the financial interests of property owners as the primary driver of restrictive regulation in the US. The interests of housing developers are the counterweight that prevent regulation from choking off all new construction (Ouasbaa et al., 2025).

In the tradition of Glaeser and Gyourko (2002), it will be helpful to interpret the equilibrium markup of floorspace prices p_n over marginal construction costs $c_n(h_n)$ as a sufficient statistic for the tightness of housing supply regulation. A high markup implies the existence of a barrier to entry in the housing market. The work of D'Amico et al., 2024 suggests housing developers are quite competitive. For instance, they document that 40% of employment in single family home construction in the US is in firms of fewer than five employees. This implies that there are not large fixed costs associated with entering this market. Given an absence of economic barriers to entry, a natural conclusion is that markups reflect regulatory barriers to entry.

In light of this, I rearrange the municipality's optimality condition (1) for neighbourhood n to yield a markup formula:

$$\underbrace{\frac{p_n - c_n(h_n)}{p_n}}_{\text{Markup}} = \underbrace{\frac{1 - \xi_n}{\xi_n}}_{\text{Ability to regulate}} \cdot \underbrace{\sum_{n' \in \mathcal{N}_m} \left| \frac{d \ln p_{n'}}{d \ln h_n} \right| \frac{p_{n'} h_{n'}}{p_n h_n}}_{\text{Incentive to regulate}} \tag{2}$$

This formula states that the equilibrium tightness of housing supply regulation can be written as the product of two terms. The first, labelled the 'ability to regulate', represents the relative influence of incumbents over developers in the municipality's decision-making. I take it as a fixed institutional feature of the neighbourhood. The second term, labelled the 'incentive to regulate', measures the size of the pecuniary externalities that new construction imposes on property owners throughout municipality m. It is a sum of inverse cross-price demand elasticities, weighted by the value of the housing stock in each neighbourhood. The larger are these elasticities, the more prices fall when new housing is built, and therefore the larger is the incentive for municipalities to tighten regulation in favour of property owners.

2.1.1 Discussion of Approaches to Representing Regulation

Although most other papers treat regulation as either non-existent or exogenous, there have been nonetheless several approaches to representing it in a model. There are those that focus on a specific aspect of regulation and represent it faithfully within their model. The clearest examples of this are in the quantitative literature on minimum lot size regulations (Kulka, 2019; Song, 2021; Macek, 2024), where these regulation are taken to be (potentially flexible) constraints on the amount of housing that a given household can consume. This type of approach has the advantage of being verifiably close to reality. However, it typically also has very demanding data requirements, such as property-level transaction databases provided by aggregators like CoreLogic, or digitised zoning maps from individual municipalities.

My approach falls into the other broad category (e.g. Glaeser and Gyourko (2002), Herkenhoff et al. (2018), Babalievsky et al. (2024), and Duranton and Puga (2023)) that aims to capture the total effect of regulation in a low dimensional, theoretically motivated way. In these models, regulation is typically a wedge between the equilibrium outcome and a certain frictionless benchmark. In my case, this takes the form of the markup of floorspace prices above marginal costs. If there were no motive for municipalities to regulate the housing supply ($\xi_n = 1$ for all n), then prices would be equal to marginal costs. This has the converse advantages and disadvantages. While it does lose immediate verifiability, it does allow the researcher to study regulation with minimal data requirements. It also plays the role of a 'sufficient statistic' for any number of specific types of regulation, as long as they present a barrier to entry to developers.

2.2 Housing Demand

There is a continuum $\mathcal{L}$ of households with measure L. Aside from the city, there is also a fixed outside option, assumed to be the rest of the country. Households l are differentiated by three characteristics: 1) a location-independent exogenous skill type $\mathfrak{s}(l) \in \mathcal{S}$; 2) taste shocks for the city, $\mathfrak{y}_{\mathcal{N}}(l)$, and the outside option, $\mathfrak{y}_{\emptyset}(l)$; and 3) residence- and workplace-specific taste shocks $(\mathfrak{z}_{nk}(l))_{n,k\in\mathcal{N}}$.

Households first choose whether to live in the city or to choose the outside option, which gives them skill-specific utility $\bar{V}_s$. Households then choose a neighbourhood n to live in, and a neighbourhood k to work in. In their residential neighbourhood, households consume housing floorspace at price p_n, as well as a tradable numeraire good. They enjoy exogenous amenities b_{ns} specific to skill type s, as well as endogenous amenities η_n which I proxy by the neighbourhood's average income (Lee and Lin, 2018; Macek, 2024). They earn a wage w_{ks} in their workplace neighbourhood, and suffer a utility cost as a function of commuting time τ_{nk}.

2.2.1 Housing and Tradable Good Consumption

I first analyse the problem of a household of skill type s which has already chosen its residence neighbourhood n and workplace neighbourhood k. It earns a wage w_{ks}, and uses this to buy housing floorspace at price p_n[2] and a tradable good whose price is normalised to one without loss of generality. Given housing consumption h and tradable good consumption c, households maximise a non-homothetic CES consumption aggregator defined implicitly as follows (see Comin et al. (2021) and section A.1 for more details):

$$\begin{aligned} \max_{c,h} \quad & C(c,h) \\ \text{s.t.} \quad & C(c,h) = \left(c^{\frac{\sigma-1}{\sigma}} + C(c,h)^{\frac{\nu}{\sigma}} h^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \\ & c + p_n h = w_{ks} \end{aligned} \tag{3}$$

Many quantitative urban models assume a Cobb-Douglas aggregator instead (Ahlfeldt et al., 2015; Monte et al., 2018; Severen, 2018; Bordeu, 2023). This buys tractability, since it means price elasticities of demand for a location are constant.[3] As stated in section

[2]In a static model, there is no distinction between renting and buying housing. In both cases, households are simply paying for housing consumption. The more important assumption is that households are not being rebated the rents from owners of housing. If they were, it could potentially alter their location decisions.

[3]One may allow for different Cobb-Douglas shares among different skill groups, but this does not allow for different elasticities for rich and poor households within the same group.

1, it is crucial that my model match the empirical reality that rich households are less price elastic than poor households. Non-homothetic CES preferences can achieve this. Moreover, there is robust empirical evidence that housing demand is non-homothetic (see, for example, Figure B.2). That is, rich households spend a smaller share of their income on housing.

The parameter ν controls the shape of the Engel curve. If $\nu < 0$, as will be calibrated in the quantification of the model, then households with higher overall consumption will place less importance on housing relative to tradable good consumption. Naturally, this means they will spend a smaller share of their income on it. Given that richer households are able to consume more in general, this implies that housing expenditure shares decline in income.

As derived in section A.1, the household's real wage (i.e. the value function of the maximisation problem) is defined by

$$\frac{w_{ks}}{P(p_n, w_{ks})} \tag{4}$$

where

$$P(p_n, w_{ks}) = \left(1 + \left(\frac{w_{ks}}{P(p_n, w_{ks})}\right)^{\nu} p_n^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \tag{5}$$

is the non-homothetic CES price index in neighbourhood n and workplace k for skill type s. Notice that, unlike a homothetic CES price index, it depends on the wage as well as the price of housing. It is also defined implicitly.

2.3 Residence and Workplace Choice

Anticipating its optimal choice of housing and tradable good consumption, the indirect utility from living in n and working in k for household l is the following:

$$\begin{aligned} \mathfrak{v}_{nk}(l) &= v_{nk\mathfrak{s}(l)} \cdot \mathfrak{z}_{nk}(l) \cdot \mathfrak{y}_{\mathcal{N}}(l) \\ \text{where} \qquad v_{nk\mathfrak{s}(l)} &= \frac{w_{k\mathfrak{s}(l)}}{P(p_n, w_{k\mathfrak{s}(l)})} \cdot e^{-\kappa\tau_{nk}} \cdot \eta_n^{\psi_{\mathfrak{s}(l)}} \cdot b_{n\mathfrak{s}(l)} \end{aligned} \tag{6}$$

Indirect utility is the product three main terms. The first is a common component, $v_{nk\mathfrak{s}(l)}$, which depends on the household's skill type $\mathfrak{s}(l)$. The second two are taste shocks for the residence-workplace pair (n, k), $\mathfrak{z}_{nk}(l)$, and the city, $\mathfrak{y}_{\mathcal{N}}(l)$.

The common component of indirect utility, $v_{nk\mathfrak{s}(l)}$, is itself the product of several terms. The first term is the real wage, as defined in subsection 2.2.1. The second term is a utility shifter which exponentially decays with the commuting time τ_{nk} between n and k. The speed of this decay is controlled by parameter κ. The final two terms, η_n and $b_{n\mathfrak{f}(l)}$, are the endogenous and exogenous amenities respectively. The strength of the endogenous amenity spillovers is skill-type specific, to accommodate the common result that college educated and richer households often are estimated to have higher preference for endogenous amenities than non-college educated and poorer households (Diamond, 2016; Macek, 2024). Conditional on living in the city, household l's location choice problem is therefore

$$\max_{n,k\in\mathcal{N}} \quad v_{nk\mathfrak{s}(l)} \cdot \mathfrak{z}_{nk}(l) \tag{7}$$

If the residence and workplace taste shocks $\mathfrak{z}_{nk}$ are distributed independently among households according to the following Fréchet distribution

$$\mathbb{P}(\mathfrak{z}_{nk}(l) \leq z) = e^{-z^{-\gamma}} \qquad \gamma > 0 \tag{8}$$

then the demand among skill type s for living in n and working in k is

$$\ell_{nks} = \frac{v_{nks}^{\gamma}}{V_s^{\gamma}} \cdot L_s \tag{9}$$

where L_s is the number of households of skill type s that choose to live in the city, and V_s is the average indirect utility among those households. From the properties of the Fréchet distribution, it takes the following expression:

$$V_s = \left(\sum_{n,k\in\mathcal{N}} v_{nks}^{\gamma} \right)^{\frac{1}{\gamma}} \tag{10}$$

This permits writing the endogenous amenities explicitly as

$$\eta_n = \frac{\sum_{k\in\mathcal{N},s\in\mathcal{S}} \ell_{nks} \cdot w_{ks}}{\sum_{k\in\mathcal{N},s\in\mathcal{S}} \ell_{nks}} \tag{11}$$

Notice that equations (6), (9) and (11) define a fixed point system. Households' location decisions depend on the average income in each neighbourhood, and the average income in each neighbourhood in turn depends on households' location decisions. This will be relevant in the discussion of inverting and solving the model numerically.

Equation (9) also allows one to obtain other important expressions. Summing across residence locations yields the labour supply to each workplace location:

$$\ell_{ks} = \frac{\sum_{n\in\mathcal{N}} v_{nks}^{\gamma}}{V_s^{\gamma}} \cdot L_s \tag{12}$$

Equation (9) can instead be combined with the expression for non-homothetic CES housing demand, derived in section A.1, to get the total housing demand in neighbourhood n:

$$h_n = \sum_{k \in \mathcal{N}, s \in \mathcal{S}} \left(\frac{p_n}{P(w_{ks}, p_n)} \right)^{-\sigma} \left(\frac{w_{ks}}{P(w_{ks}, p_n)} \right)^{1+\nu} \ell_{nks} \tag{13}$$

2.3.1 Discussion of the Price Elasticity of Housing Demand

As shown in equation (2), the elasticities of housing demand are key for determining regulation in equilibrium. There are two key features of this model of household behaviour which lead to rich households having lower demand elasticities. The first is the non-homotheticity of preferences over housing and tradable goods. While this does not always map into lower elasticities of demand for the rich, in this particular formulation it does. Crucially, with the parameters I use for the non-homothetic CES aggregator (discussed in section 3.3.1), the price index P has the following property:

Lemma 1. *If* $\nu < 0$ *and* $\sigma \in (0, 1)$,

$$\frac{\partial^2 \ln P}{\partial \ln p \partial \ln w} < 0 \tag{14}$$

Proof: See appendix A.2.

That is, price increases reduce the indirect utility of rich households proportionally less than that of poor households. Therefore, prices are less important for their location choices, and they are less price elastic as a result.

The second feature of the model which leads to differential price elasticities is the skill-specificity of the endogenous amenity spillovers. In the baseline calibration (discussed again in section 3.3.1), high-skill households value the presence of the rich more than low-skill households do. For a given price rise in a neighbourhood, this means that high-skill households will be 'compensated' by the resulting exit of poor households more than low-skill households. This provides an additional reason for price increases to be less important to them, and therefore for their housing demand elasticities to be lower.

This second point is particularly important because it highlights that my model captures a common 'alternative' rationale for regulation that is proposed in the literature; namely that its purpose is directly to keep out households of a particular socioeconomic or demographic group. In my model, these effects are present, but they all operate through the elasticities of demand for housing.

2.4 City Choice

The first choice that households have to make is whether to live in the city, or to choose a fixed outside option. One should think of this as the rest of the country in which the city is embedded, and that it is large enough not to be affected by population flows in and out of the city. This outside option gives utility $\bar{V}_s$ to households of skill type s. Households have a taste shock for the city $\mathfrak{y}_{\mathcal{N}}$, and one for the outside option $\mathfrak{y}_{\emptyset}$. Given that they do not yet know the realisation of $\mathfrak{z}_{nk}$, they choose based on their expected utility from living in the city. Therefore, their problem is

$$\max \quad \left\{ V_{\mathfrak{s}(l)} \cdot \mathfrak{y}_{\mathcal{N}}(l), \bar{V}_{\mathfrak{s}(l)} \cdot \mathfrak{y}_{\emptyset}(l) \right\} \tag{15}$$

If these taste shocks are independently Fréchet distributed with shape parameter ζ, then demand for living in the city takes the following form:

$$L_s = \frac{V_s^{\zeta}}{V_s^{\zeta} + \bar{V}_s^{\zeta}} \cdot \bar{L}_s \tag{16}$$

where $\bar{L}_s$ is the exogenous mass of households of skill type s that live in the country.

2.5 Labour Demand

In each neighbourhood k there is a price-taking representative firm that hires labour of different skill types and produces the tradable good. It uses the following technology, which has a constant elasticity of substitution across skill types:

$$q_k = \left(\sum_{s \in \mathcal{S}} a_{ks}^{\frac{1}{\rho}} \ell_{ks}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1} \cdot \delta} \tag{17}$$

The parameter δ controls the returns to scale in the production function. It can be thought of as playing the role of a fixed factor in production, such as land, which is not explicitly modelled here. Each firm takes wages for each skill type as given, and therefore inverse labour demand for type s is given by:

$$w_{ks} = \delta \cdot a_{ks}^{\frac{1}{\rho}} \cdot q_k^{\frac{\rho(\delta-1)+1}{\delta\rho}} \cdot \ell_{ks}^{-\frac{1}{\rho}} \tag{18}$$

2.6 Equilibrium

Definition 1 (Equilibrium)**.** An equilibrium is a tuple $(\boldsymbol{w}, \boldsymbol{p}, \boldsymbol{\eta}, \boldsymbol{h}, \boldsymbol{\ell})$ such that for all $n, k \in \mathcal{N}$ and $s \in \mathcal{S}$ such that local price indices are defined by (A.16); location choices are defined by (6), (9), (11), (10), and (16); housing demand defined by (13) is equal to the housing supply defined implicitly by each municipality's optimality conditions (1); labour demand, given by (17) and (18), is equal to labour supply (12).

One should think of this model as qualitatively being a series of housing markets and labour markets. These markets are cleared by arrays of floorspace prices $\boldsymbol{p}$ and skill type specific wages $\boldsymbol{w}$, and are connected by commuting flows. The key equations for the housing market are (13) and (1), while the key equations for the labour market are (18) and (12).

2.7 A Simple Example of the Key Mechanism

In this section, I outline a simple version of the model which allows for a clean exposition the interaction between income segregation and endogenous regulation. This example is also directly analogous to how I later quantify the effects of endogenous regulation in my counterfactual results.

Equation (2) decomposes markups into an institutional and economic term. Consider the special case where $\xi_n = \frac{1}{2}$ for some neighbourhood n; that is, property owners and housing developers have equal political power. Moreover, suppose that municipality m consists of only neighbourhood n, and that the city is large enough so that n is negligible in size relative to it. This permits analysis of n in partial equilibrium, since outcomes in the rest of the city can be taken as given.

Dropping subscripts for simplicity, in this case equation (2) collapses to a familiar markup formula for an oligopolist:

$$\frac{p - \bar{c}}{p} = \frac{1}{\left|\frac{d \ln h}{d \ln p}\right|} \tag{19}$$

That is, the markup of price above the marginal construction cost is equal to the inverse of the demand elasticity for housing. The municipality behaves as if it were an oligopolist maximising total profits from housing production. Crucially, in neighbourhoods with inelastic demand, it will 'exploit' that demand with a high markup.

To understand how the municipality will adjust regulation in response to the income composition of the neighbourhood, one must specify the demand side of the housing market. For simplicity's sake, I assume that each household consumes one unit of housing,[4] and households have the following skill-specific housing demand curves:

$$h_L(p) = p^{-\varepsilon_L} \qquad h_H(p, w) = w \cdot p^{-\varepsilon_H} \tag{20}$$

$$\varepsilon_H < \varepsilon_L \qquad w > 1 \tag{21}$$

The shifter w represents the high-skill wage; low-skill wages are implicitly normalised to one. Crucially, high-skill households are less price elastic than low-skill households. Conditional on the share of high-skill households in the neighbourhood, which I denote with π, the local elasticity of housing demand is the following:

$$\left|\frac{d \ln h}{d \ln p}\right| = \pi \cdot \varepsilon_H + (1 - \pi) \cdot \varepsilon_L \tag{22}$$

It is a convex combination of the high-skill and low-skill elasticities of demand, weighted by the high-skill share π. In the full quantitative version of the model, this elasticity is far more complicated. It involves spatial interactions, and the elasticities of different groups are not constant, but nonetheless the key logic remains: richer neighbourhoods are less price elastic, all else equal.

Denoting the markup as $\mu \equiv \frac{p-\bar{c}}{p}$, the housing market equilibrium can be defined by the following two equations in π and μ:

Household behaviour:
$$\pi = f_\pi(\mu, w) \equiv \frac{h_H\left(\frac{\bar{c}}{1-\mu}, w\right)}{h_L\left(\frac{\bar{c}}{1-\mu}\right) + h_H\left(\frac{\bar{c}}{1-\mu}, w\right)} \tag{23}$$

Municipality behaviour:
$$\mu = f_\mu(\pi) \equiv \frac{1}{\pi \cdot \varepsilon_H + (1 - \pi) \cdot \varepsilon_L} \tag{24}$$

The first equation defines the how the income composition of the neighbourhood responds to the tightness of regulation (namely, the markup). The second equation, in turn, describes how the municipality chooses regulation (the markup) in response to the income composition of the neighbourhood. This system therefore summarises the feedback loop between neighbourhood income and housing supply regulation. High markups push out poor households and make the neighbourhood richer (equation 23), and municipalities choose tighter regulations for richer neighbourhoods (equation 24).

This feedback loop can be illustrated in a graphical example. Suppose that the

[4] This is not necessary for the mechanism, but it makes analysis of the simple case cleaner.

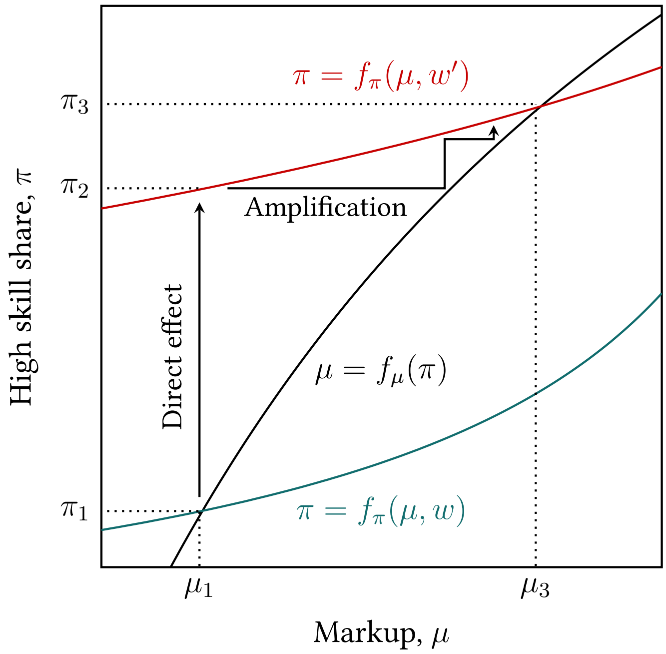

Figure 1: How regulation adjusts after an increase in the college wage premium

high-skill income is initially equal to some $w > 1$. Figure 1 represents this initial equilibrium with the intersection of the green line, which represents the households' response to the municipality's markup (equation 23), and the black line, which represents the municipality's response to the households' location choices (equation 24). The intersection of these lines defines the initial equilibrium markup, μ_1, and high-skill share, π_1.

Suppose that the college wage premium increases; that is, the wage of high-skill households increases from w to w'. Figure 1 shows that this has the effect of raising the high-skill share for any given markup; the green line becomes the red line. The result of this can be decomposed into two effects. The first, which I call the 'direct effect', is what happens if the municipality does not adjust regulation in response to the changing income distribution of the neighbourhood. In this simple case, that means it behaves as if the high-skill share stays fixed at π_1. As such, it keeps its markup constant at $\mu_1 = f_\mu(\pi_1)$, and the high-skill share rises to π_2.

What happens if the municipality then starts responding optimally to the local income distribution? Given the high-skill share of π_2, the municipality's optimal choice is to charge a higher markup. Households re-sort in response to this higher markup, and the high-skill share increases further. This cycle continues, represented by the zig-zag pattern in the arrow in Figure 1, until households' and municipalities' choices are both consistent with one another. The high-skill share ends up even higher than before, at π_3.

I call this feedback loop the 'amplification effect'.

3 Quantification

In this section I outline how I bring the model to real world data. There are two main types of unobserved quantities that are needed to conduct counterfactual solutions. The first are the location fundamentals, $\left(\boldsymbol{\xi}, \boldsymbol{a}, \boldsymbol{b}, \bar{\boldsymbol{V}}\right)$. These are location-specific terms, and are usually exactly identified from the data. Hence, the process of recovering them is often referred to as 'inversion'. The second group are the parameters, $(\sigma, \nu, \gamma, \zeta, \psi_L, \psi_H, \rho, \delta)$. These do not depend on location. Instead, they reflect general properties of preferences and technology.

In section 3.1, I describe the data sources that I use. In section 3.2, I detail the process for inverting the observed equilibrium to recover the location fundamentals, conditional on the set of parameters. In section 3.3, I then describe how I calibrate and estimate these parameters. Finally, in 3.4 I show how the model is able to match key first-order facts from the data.

3.1 Data

3.1.1 RSMeans Construction Cost Books

A key part of my analysis relies on measuring the markup of housing prices above marginal construction costs. Data on housing prices is relatively common; data on construction costs is not. I use a pair of books on construction costs published by RSMeans in 2017,[5] a large American industrial data aggregator. These books are freely available in many US public libraries. In large part I follow the approach of Glaeser and Gyourko (2002), who construct markups from this dataset at the metropolitan area level. My key departure from their work is that I construct markups instead at the census tract level, using RSMeans' breakdowns of construction costs by structure type, and another dataset on the mix of structure types in each census tract.

I start by digitising two groups of tables from these books using GPT-4.1's image recognition API. The first group of tables details the average marginal cost of building a square foot of housing with different characteristics. The most important of these

[5] In fact, one of the books was only available from 2019. I deflate all estimates from this book by a nationwide construction cost index.

characteristics for the purposes of this paper is the type of structure that a housing unit belongs to. The tables provide average construction costs for four basic types: single family homes; one to three storey apartment buildings; four to seven storey apartment buildings; and eight to twenty-four storey apartment buildings. As mentioned, the ultimate aim with this data is to combine these construction costs with other data on the observed mix of structure types in each census tract (described below). In order to make these two datasets compatible, I collapse the two largest categories in the RSMeans data into 'large apartment buildings', and refer to the second category as 'medium size apartment buildings'. The tables also provide cost breakdowns by other building characteristics, such as overall quality and building materials. Following Glaeser and Gyourko (2002), I use only cost estimates for 'economy' housing units (the lowest quality type), and average across other characteristics.

The second group of tables provides location-specific adjustment factors to take into account differences in terrain suitability and local labour costs. These 'locations' are groups of three-digit zip codes, covering all major urban areas. In total, there are 862 locations in the table. I multiply the marginal construction costs from the first table by these adjustment factors. The next step is to allocate these location and structure type specific marginal costs to census tracts, for which I use the American Community Survey.

3.1.2 American Community Survey

I use the American Community Survey (ACS) tabulated at the census tract level (accessed from NHGIS). Census tracts are drawn to contain roughly 4,000 people each, and I follow other papers in this literature by treating them as the empirical equivalent of the 'neighbourhood' in my model (Almagro et al., 2024; Couture et al., 2024; Owens III et al., 2020). There are roughly 75,000 census tracts in the United States, and after dropping missing observations I am left with a working sample of 56,294 census tracts covering 851 metropolitan and micropolitan areas.[6] I treat these metropolitan and micropolitan areas as the equivalent of a 'city' in my model, since they are large enough to be considered a self-contained labour market. Most of my analysis is conducted using the 2015-2019 ACS sample, which I will refer to for convenience's sake as 2017 (the central year). I choose this sample because it is the last ACS sample to be entirely unaffected by the Covid-19 pandemic and its associated changes to location choices.

To map the RSMeans construction cost estimates to census tracts, I require data on the mix of structure types in each tract. While the ACS does have this data, it does not

[6]Officially known as Core Based Statistical Areas.

always define structure types in the same way as RSMeans. For single family homes this is not the case; they both define them the same way. However, for apartment buildings RSMeans divides structure types by their number of storeys, whereas the ACS divides structure types by the number of housing units in each structure. Therefore I need to take a stance on how to define the boundary between 'medium size apartment buildings' (one to three storeys) and 'large apartment buildings' (over four storeys) using the ACS data. Luckily, RSMeans provides plausible ranges of the total floorspace in each structure type, which gives a rough estimate of the expected number of housing units. Using this, I assign apartment buildings in the ACS with over fifty housing units to be 'large', and all other apartment buildings to be 'medium size'.

While I believe this is a good approximation, I also do not consider it to be a major concern. The primary determinant of a structure type's construction cost per square foot is whether it is a single family home or not. Apartment buildings are simply much more expensive to build, regardless of whether they are large or small. Therefore, the most important dimension of cost variation is whether developers are building single family homes or not, which is precisely measured in the ACS. Moreover, large apartment buildings are simply very rare in the urban US.

Armed with data on construction costs by structure type, and structure type shares by census tract, I describe in section 3.2.4 how I estimate the overall marginal construction cost in each neighbourhood. To compute markups, I also need data on housing prices per square foot in each census tract. Given that square feet of housing are not reported in the census, I follow a very similar procedure to Hoelzlein (2023): I multiply median housing expenditures by the number of housing units in the neighbourhood, and divide by the total number of rooms. This gives a measure of housing expenditure per room. The remaining step is then to rescale this to be housing expenditure per square foot. I do this by computing the median model-implied housing demand among all households, and dividing this by the US-wide median size of a housing unit, measured in square feet, from the American Housing Survey.[7] I then divide prices by this ratio to obtain the model-implied price per square foot.

3.1.3 Other Data Sources

I augment the ACS with data from the LEHD Origin-Destination Employment Statistics to obtain residence- and workplace-specific census tract employment numbers for college and non-college educated workers in 2017. I treat these classes of workers as the empirical

[7]In future work, I plan to use MSA-specific medians from the American Housing Survey.

equivalents of the skill types in the model, in line with other work (Diamond, 2016; Giannone, 2022; Hoelzlein, 2023). I also use the Urban Institute's bilateral driving time matrices at the census tract level. These two datasets, along with the ACS variables, allow me to recover unobserved quantities (developer power, wages, amenities, and productivities) from the observed equilibrium. I outline this process in section 3.2.

Finally, for the sake of validating my model-implied measures of regulation, I use the 2018 update of the Wharton Residential Land Use Regulation Index (Gyourko et al., 2008; Gyourko et al., 2021). This is a survey-based index measuring the tightness of housing supply regulation in around 2,500 zoning jurisdictions in the US, focusing on the suburbs of large metropolitan areas. Since it does not cover entire cities, nor does it measure regulatory tightness for different census tracts within municipalities, it cannot be used as a data input to the model. Nonetheless, it can be used as a validation of my model-implied measures of regulation.

3.2 Model Inversion

For now I treat the location-independent parameters in the model as known. In section 3.3, I detail how I calibrate and estimate them. In this section, I outline the procedure for inverting the model's location fundamentals. In this case, they are: amenities $\{b_{ns}\}$; productivities $\{a_{ks}\}$; indirect utilities of the outside option $\{\bar{V}_s\}$; and developer power $\{\xi_n\}$. There will also be fundamentals derived from the form I choose for the marginal construction cost functions, $\{c_n\}$, that I introduce below.

3.2.1 Wages

Inverting the model requires data on equilibrium outcomes. In this case, skill-specific wages by workplace neighbourhood are unobserved, and I have to recover them from data on other equilibrium outcomes. This is a common procedure in quantitative urban models (see e.g. Ahlfeldt et al. (2015) and Tsivanidis (2023)), since data on wages at this level typically do not exist. Following Ahlfeldt et al. (2015), I rewrite the labour supply equation (12) so that it conditions on the number of households by residential location and skill type.

$$\ell_{ks} = \sum_{n \in \mathcal{N}} \frac{\left(\frac{w_{ks}}{P(w_{ks}, p_n)} e^{-\kappa \tau_{nk}}\right)^{\gamma}}{\sum_{k' \in \mathcal{N}} \left(\frac{w_{k's}}{P(w_{k's}, p_n)} e^{-\kappa \tau_{nk'}}\right)^{\gamma}} \cdot \ell_{ns} \tag{25}$$

Given data on households by residential location $\{\ell_{ns}\}$, households by workplace location $\{\ell_{ks}\}$, and commuting times $\{\tau_{nk}\}$, this equation can be solved via a standard tâtonnement procedure. An important difference with Ahlfeldt et al. (2015) is that here, due to the non-homotheticity, floorspace prices do not cancel out of the equation. This may seem innocuous, but it strengthens the requirements for knowledge of parameters before passing this step of the inversion process. In the homothetic case, one only needs to know the product $\gamma \cdot \kappa$ in order to back out a transformation of the wage, w_{ks}^{γ}. This is particularly convenient because this product is readily estimatable in a standard 'commuting gravity' regression using publicly available travel survey data. In contrast, with non-homotheticity one needs to take a stance on the 'untransformed' wage in each step of the iteration in order to be able to compute the price index $P(w_{ks}, p_n)$. As such, the econometrician must take a stance on both γ and κ.

3.2.2 Amenities and Productivities

Given wages $\{w_{ks}\}$, the next step is to recover exogenous amenities $\{b_{ns}\}$ and productivities $\{a_{ks}\}$. To recover amenities, one can sum equation (9) across workplaces to obtain an expression for the total number of workers of each skill type living in neighbourhood n. Given that amenities are only identified up to a multiplicative constant, I normalise the amenities in some neighbourhood $0 \in \mathcal{N}$ to one for both skill types. As a result, amenities can be written in closed form as

$$b_{ns} = \left(\frac{\ell_{ns}}{\ell_{0s}}\right)^{\frac{1}{\gamma}} \left(\frac{\left[\sum_{k \in \mathcal{N}} \left(\frac{w_{ks}}{P(p_n, w_{ks})} \cdot e^{-\kappa \tau_{nk}}\right)^{\gamma}\right] \cdot \eta_n^{\psi_s \gamma}}{\left[\sum_{k \in \mathcal{N}} \left(\frac{w_{ks}}{P(p_0, w_{ks})} \cdot e^{-\kappa \tau_{0k}}\right)^{\gamma}\right] \cdot \eta_0^{\psi_s \gamma}}\right)^{-\frac{1}{\gamma}} \tag{26}$$

The intuition for this equation comes from the idea of spatial equilibrium. Neighbourhoods with many households despite poor observed characteristics (low wages, high prices, and low endogenous amenities) must have high exogenous amenities. If not, those neighbourhoods would not be able to attract so many households.

Note from equation (13) that the quantity of housing h_n in neighbourhood n depends on the price of housing p_n, wage $\{w_{ks}\}$, and households' commuting choices $\{\ell_{nks}\}$. In turn, commuting choices in equation (9) depend on wages, prices, endogenous amenities η_n, exogenous amenities $\{b_{ns}\}$, and the mass of workers of each skill type living in the city, $\{L_s\}$. Recovering exogenous amenities therefore allows me to compute the model-implied quantity of housing in each neighbourhood. This will be useful in recovering developer power from the data, as described in section 3.2.5.

Unlike exogenous amenities, it is not possible to invert the productivity terms in closed form, but the procedure for solving for them is nonetheless computationally straightforward. Writing equations (17) and (18) together, we have the inverse labour demand curve for each skill type as a function of the productivity terms.

$$w_{ks} = \delta \cdot a_{ks}^{\frac{1}{\rho}} \cdot \left(\sum_{s \in \mathcal{S}} a_{ks}^{\frac{1}{\rho}} \ell_{ks}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho(\delta-1)+1}{\rho-1}} \cdot \ell_{ks}^{-\frac{1}{\rho}} \tag{27}$$

Given knowledge of wages $\{w_{ks}\}$ and labour supply $\{\ell_{ks}\}$, this can be solved for via a tâtonnement procedure. Since it does not involve cross-neighbourhood interactions, it is very fast to solve.

3.2.3 Utility of the Outside Option

Once wages $\{w_{ks}\}$ and amenities $\{b_{ns}\}$ have been solved for, from equation (15) I am able to compute the average indirect utility of households who live in the city, $\{V_s\}$. As a result, this allows me to invert the implied indirect utility of the outside option for each skill type.

$$\bar{V}_s = V_s \left(\frac{\bar{L}_s - L_s}{L_s} \right)^{\frac{1}{\zeta}} \tag{28}$$

Given that I am framing the outside option as the rest of the country, I take $\bar{L}_s$ to be the total number of households (workers) of type s in the country from the LODES data.

3.2.4 Housing Cost Fundamentals

I have so far left the other key production technology in the model, the marginal construction cost function $c_n(h_n)$, unspecified. I do this because none of the model depends substantively on the functional form of c_n. However, to conduct counterfactual exercises one must take a stance on it. Here, I outline a simple microfoundation whose fundamentals I can recover from the publicly available data I use. The aim is to capture the congestion forces that arise as a neighbourhood becomes denser while retaining minimal data requirements.

I assume that all floorspace has to be built as part of a structure, and there are three structure types. Let $\mathcal{I} = \{\text{single family}, \text{medium size multifamily}, \text{large multifamily}\}$ be the set of these structure types. There is a constant, neighbourhood-specific marginal cost of building a unit of floorspace for each structure type. I denote this marginal cost

$\aleph_{ni}$ for structure type i in neighbourhood n.

For each amount of built floorspace h_n, there is a neighbourhood-specific exogenous share $\phi_{ni}(h_n)$ that must be devoted to each structure type i. This captures the scarcity of land: on empty land, it is easy to build only single family housing. As the land fills up, developers must start building denser forms of housing. The total cost function for neighbourhood n, which I denote by g_n, is therefore

$$g_n(h_n) = h_n \cdot \sum_{i \in \mathcal{I}} \phi_{ni}(h_n) \cdot \aleph_{ni} \tag{29}$$

Furthermore, I assume that the structure type shares ϕ_{ni} take the following logistic form:

$$\phi_{ni}(h_n) = \frac{\beth_{ni} \cdot \left(\frac{h_n}{\lambda_n}\right)^{\daleth_i}}{\sum_j \beth_{nj} \cdot \left(\frac{h_n}{\lambda_n}\right)^{\daleth_j}} \tag{30}$$

I denote the total land area of neighbourhood n as λ_n, and therefore ratio $\frac{h_n}{\lambda_n}$ as housing density. The neighbourhood-specific shifters $\beth_{ni}$ capture idiosyncratic variation in structure type shares that is unrelated to this housing density. For instance, for historical reasons older neighbourhoods may have a larger share of multifamily structures than similarly dense newly built neighbourhoods.

The elasticities $\daleth_i$ control the speed with which structure type i increases or decreases its share as the neighbourhood gets denser. These parameters are key for counterfactual experiments in which the analyst needs to know how marginal construction costs will respond to changes in the quantity of housing. Differentiating with respect to h_n, marginal construction costs are given by

$$c_n(h_n) = \sum_{i \in \mathcal{I}} \aleph_{ni} \phi_{ni}(h_n) \left(1 + \daleth_i - \sum_{j \in \mathcal{I}} \phi_{nj}(h_n) \daleth_j\right) \tag{31}$$

Since this is a novel formulation of the construction cost function, I estimate its parameters $\{\daleth_i\}$ myself. I describe this procedure in section 3.3.2. Here I take those parameters as given, and outline the procedure for recovering the location fundamentals $\{\aleph_{ni}\}$ and $\{\beth_{ni}\}$.

The first step is to note that the structure-type shares $\{\phi_{ni}(h_n)\}$ are unique in $\{\daleth_i\}$ up to an additive constant. Therefore, one can additively normalise $\daleth_{\text{single family}}$, the parameter that determines the share of single family housing as a function of housing density, to zero. The remaining parameters, $\daleth_{\text{medium size multifamily}}$ and $\daleth_{\text{large multifamily}}$, can subsequently be interpreted as being relative to $\daleth_{\text{single family}}$. Similarly, the structure-type shifters $\{\beth_{ni}\}$ can be normalised multiplicatively so that $\beth_{n,\text{single family}} = 1$ for all neighbourhoods n. As a result, one can recover the shifters from the following equation:

$$\beth_{ni} = \frac{\phi_{ni}}{\phi_{n1}} \left(\frac{h_n}{\lambda_n} \right)^{-\daleth_i} \tag{32}$$

3.2.5 Developer Power

The final and most difficult step of inverting my model from an observed equilibrium is the recovery of developer power $\{\xi_n\}$. From equation (1) this can be written as

$$\xi_n = \frac{\sum_{n' \in \mathcal{N}_m} \frac{dp_{n'}}{dh_n} h_{n'}}{\sum_{n' \in \mathcal{N}_m} \frac{dp_{n'}}{dh_n} h_{n'} - (p_n - c_n(h_n))} \tag{33}$$

Prices p_n are observed, and both quantities of housing h_n and marginal costs $c_n(h_n)$ are easily computable from the data and previously inverted fundamentals. The only part that is not readily computable is the inverse cross-price derivative of the housing demand function, $\frac{dp_{n'}}{dh_n}$.

There are several reasons why this is not straightforward to compute. The first is that the inverse housing demand function is not expressible in closed form. In fact, neither is the housing demand function itself (equation 13) due to the presence of the non-homothetic CES price index, but this is fast to compute via fixed point iteration or linear interpolation. Therefore, no matter the method for computing derivatives, one must compute the Jacobian of housing demand and then invert it. The elements of the resulting matrix will be the desired inverse cross-price derivatives.

I solve for the Jacobian matrix in the following way.[8] First, I write its terms as a function of other model quantities. Given the multiplicative structure of the model, it is convenient to write out the cross-price elasticities, rather than cross-price derivatives:

$$\begin{aligned} \frac{d \ln h_n}{d \ln p_j} = \sum_{ks} \frac{h_{nks}}{h_n} \Bigg[\mathbb{1}\{n = j\} \Bigg(-\sigma \left(1 - \frac{d \ln P(w_{ks}, p_n)}{d \ln p_n} \right) \\ -(\nu + 1) \frac{d \ln P(w_{ks}, p_n)}{d \ln p_n} \Bigg) + \frac{d \ln \ell_{nks}}{d \ln p_j} \Bigg] \end{aligned} \tag{34}$$

This in turn contains two different groups of elasticities. The first, $\{\frac{d \ln P(w_{ks}, p_n)}{d \ln p_n}\}$, is easily

[8] The simplest approach would be to use a finite differences approximation. In appendix B.4 I describe why this is infeasible in my setting.

computed via finite difference methods because it does not involve cross-neighbourhood interactions. It requires evaluating P for every (n, k, s) combination twice, which can be done quickly.

The difficulty lies with the second group of elasticities, $\{\frac{d\ln \ell_{nks}}{d\ln p_j}\}$. These are the elasticities of demand for a residence-workplace pair (n, k) among skill type s with respect to the price in neighbourhood j. Unlike the first group, it is infeasible to compute them via finite differences for the same reasons discussed in B.4. As such, I differentiate (9) and (11) to write them as a function of other elasticities:

$$
\begin{aligned}
\frac{d\ln \ell_{nks}}{d\ln p_j} &= \gamma \left[\mathbb{1}\{n=j\} \left(-\frac{d\ln P(w_{ks}, p_j)}{d\ln p_j} \right) - \sum_{k'} \frac{\ell_{jk's}}{L_s} \left(-\frac{d\ln P(w_{k's}, p_j)}{d\ln p_j} \right) \right] \\
&\quad + \gamma \left[\psi_s \frac{d\ln \eta_n}{d\ln p_j} - \sum_{n'} \frac{\ell_{n's}}{L_s} \psi_s \frac{d\ln \eta_{n'}}{d\ln p_j} \right]
\end{aligned} \tag{35}
$$

$$
\frac{d\ln \eta_n}{d\ln p_j} = \sum_{ks} \frac{w_{ks} - \eta_n}{\eta_n} \frac{\ell_{nks}}{\ell_n} \frac{d\ln \ell_{nks}}{d\ln p_j} \tag{36}
$$

This defines a fixed point system of equations in $\{\frac{d\ln \ell_{nks}}{d\ln p_j}\}$ and $\{\frac{d\ln \eta_n}{d\ln p_j}\}$ (assuming that $\{\frac{d\ln P(w_{ks}, p_n)}{d\ln p_n}\}$ has already been computed). The reason that it must be written as a fixed point system, and not in closed form, is the presence of endogenous amenities. The average income in each neighbourhood is the result of households' location choices, and those choices are in turn determined by the average income in each neighbourhood. Therefore, any change in prices will change location choices. These location choices will change endogenous amenities, which in turn affect location choices again. The true elasticities must satisfy this internal consistency.

This system can be solved via fixed point iteration. However, this approach is practically difficult because of the size of the arrays. The total number of equations in this system is $|\mathcal{N}|^3 \cdot |\mathcal{S}| + |\mathcal{N}|^2$. In New York, for instance, this system has more than 128 billion equations. The process of solving these equations iteratively involves updating arrays of this size in memory, which takes up a significant amount of time.[9] As a result, solving the full system takes several hours for large metropolitan areas. This is not an issue if it only needs to be done once, as is the case when recovering $\{\xi_n\}$. However, when solving for counterfactual scenarios, this system needs to be solved many times to evaluate convergence criteria (described in section 4.1). It is therefore of great practical value to speed this computation up.

To this end, I take advantage of a simplification of the system that arises from the

[9]Even when using pre-allocated arrays.

gravity structure of commuting demand. This can be expressed in the following way:

$$\frac{d\ln \ell_{nks}}{d\ln p_j} = \frac{d\ln \ell_{nk's}}{d\ln p_j} \qquad \forall k, k' \in \mathcal{N} \quad \text{if} \quad n \neq j \tag{37}$$

That is, a large number of the elements of this array are duplicates. Conditional on the residence neighbourhood n, perturbing prices in some other neighbourhood j has the same proportional response on commuting demand to all destinations k. Paring down the arrays to include only unique elements reduces the size to scale with $O\left(|\mathcal{N}|^2\right)$, which removes memory allocation as the main bottleneck. The system is much more manageable as a result, it is possible to solve this system in under two minutes for any metropolitan area. Once the system is solved, it is straightforward to compute the full matrix of cross-price elasticities in equation (34), invert it, and obtain $\{\xi_n\}$ from equation (33).

3.3 Parameters

The model contains a number of location-independent parameters, and knowledge of these is necessary for counterfactual solutions of the model. In this section I outline how I calibrate and estimate these parameters. See table B.1 for a summary of the parameters used.

3.3.1 Household Utility Parameters

The parameters of each household's utility function are crucial in my quantitative exercise because they control the shape of the housing demand curve in each neighbourhood, as well as how the slope of this demand curve changes with the neighbourhood's income composition.

Two of these parameters are σ and ν, which come from the non-homothetic CES consumption aggregator between housing and tradable goods. In particular, σ is the elasticity of substitution across these two goods, and ν controls the slope of the Engel curve of housing expenditure. The latter is of fundamental importance for the predictions of my model. A negative value of ν implies that housing expenditure shares are declining in income, i.e. that housing demand is non-homothetic. There is substantial existing evidence for the non-homotheticity of housing demand (Finlay and Williams, 2022; Macek, 2024). I additionally provide evidence for this in Figure B.2. I use the ACS microdata sample from 2017 to plot household income against the ratio of home value to household income. As the figure shows, richer households own homes that are less valuable as a fraction of their income, indicating a lower overall expenditure burden.

In light of this, I take the parameters governing the non-homothetic CES aggregator from Finlay and Williams (2022), who estimate them from the US Panel Survey of Income Dynamics. They estimate $\sigma = 0.5$, implying that tradable goods consumption and housing are complementary. More importantly, they estimate $\nu = -0.3$. This is consistent with the reduced form pattern of housing expenditure shares declining in income.

Another important pair of parameters in the households' utility function are the endogenous amenity spillover elasticities, ψ_L and ψ_H. These determine the strength of low- and high-skill households' preferences for living in richer neighbourhoods. To the best of my knowledge, the only applicable estimates in the literature are those of Macek (2024). He estimates these elasticities for low-, medium-, and high-income groups in the US using a 'spatial BLP instrument' strategy (see also Almagro et al. (2024)), leveraging the placement of natural amenities to generate substitution across neighbourhoods and exogenously shift average incomes. Given that income and skill are very highly correlated, I take ψ_L to be the low-income estimate, and ψ_H to be the high-income estimate. Concretely, I take $\psi_L = 0.1$ and $\psi_H = 0.3$. This is consistent with findings in other papers (Diamond, 2016; Couture and Handbury, 2020) that high-skill households derive more utility from living in affluent neighbourhoods than do low-skill households.

The parameter κ governs the utility costs of commuting. I take $\kappa = 0.01$ from Ahlfeldt et al. (2015). This implies that households' indirect utility declines by roughly one percent for every additional minute spent travelling to work.

Finally, the Fréchet shape parameters γ and ζ govern the dispersion of the taste shock distributions across different choice nests: the choice between neighbourhoods within the city, and the choice between the city and the outside option. Higher values of these parameters reduce the taste shocks' dispersion, thereby increasing the relative importance of the common component of indirect utility (v_{nks} and V_s for the inner and upper nests respectively) for determining location choices. As such, these parameters are often loosely referred to as 'migration elasticities', in that they affect the elasticity of population in a location to the indirect utility of living in that location.[10]

[10] Note that these parameters are not mathematically equivalent to the stated elasticity unless there is an uncountable number of locations. This is due to the adjustment that comes from the 'multilateral resistance' term in the denominator of choice probability expressions. For instance, in the case of equation (16), the true 'migration elasticity' would be:

$$\frac{d \ln L_s}{d \ln V_s} = \zeta \cdot \left(1 - \frac{L_s}{\bar{L}_s}\right) \tag{38}$$

The migration elasticity across residential and workplace locations, γ, has been estimated in several different papers. Estimates vary substantially depending on the setting and the estimation method. The lowest estimates are around two to three, from papers like Monte et al. (2018), Severen (2018), and Tsivanidis (2023), and the highest estimates range from five to seven, from papers like Heblich et al. (2020) and Ahlfeldt et al. (2015). In the absence of a good criterion to choose between these, I take a midpoint of $\gamma = 4$. For the migration elasticity across larger locational units, I follow Desmet et al. (2018) and set $\zeta = 2$. Intuitively, this means that households are more willing to substitute across residential and workplace neighbourhoods than they are between living in the city and living elsewhere.

3.3.2 Production Parameters

The technology for production of the tradable good in each neighbourhood is governed by two parameters: ρ, the elasticity of substitution between low- and high-skill workers, and δ, the labour share in production. I choose $\rho = 1.66$, following Diamond (2016) who estimates this using the interaction of Bartik labour demand shocks and housing supply constraints. For the labour share in production, I take $\delta = 0.9$ from Desmet and Rappaport (2017).

The other remaining production parameters are those that that govern the marginal construction cost curves $\{c_n\}$. These parameters are the elasticities $\{\daleth_i\}$, which control how the shares of different structure types move with housing density. To estimate these, I rewrite equation (32) into a regression equation:

$$\ln \frac{\phi_{ni}}{\phi_{n,\text{single family}}} = \ln \bar{\beth}_i + \daleth_i \cdot \ln \frac{h_n}{\lambda_n} + \ln \tilde{\beth}_{ni} \tag{39}$$

where the structure-type shifters have been decomposed without loss of generality into a (geometric) mean component $\bar{\beth}_i$ and a residual term, $\tilde{\beth}_{ni}$. To estimate this equation, I start with ACS data on the number of housing units by structure type. Note that $\{\phi_{ni}\}$ are shares of floorspace, not shares of housing units. If housing units in different structure types differ in their average floorspace, there will be measurement error in the dependent variable of equation (39). Unfortunately, the ACS does not collect information on housing units' floorspace, so in the absence of a better measure I use this. If the size of the measurement error is uncorrelated with density, then the measurement error is classical and OLS is still able to estimate the coefficient consistently. This would be

rather than simply ζ.

violated, for instance, if houses were systematically smaller relative to apartments in denser neighbourhoods. It is not obvious that this should be the case.

To estimate this equation, I use data from the ACS on structure type shares by census tract throughout the US, as well as data on land areas λ_n and model implied housing quantities h_n. Pooling data on structure-type shares from the ACS and housing density from multiple metropolitan areas, this equation can be estimated with OLS if the regressor, $\ln \frac{h_n}{\lambda_n}$ is uncorrelated with the error term, $\ln \tilde{\daleth}_{ni}$. Intuitively, this condition means that unobserved shifters of the relative shares of medium and large apartment buildings relative to single family homes must be uncorrelated with density.

Table 1 shows the results of this regression. Given that there are some neighbourhoods without any apartment buildings, I take the inverse hyperbolic sine of the dependent variables rather than the natural logarithm. The results confirm the intuition that as a neighbourhood becomes denser, it must start building apartment buildings instead of single family homes.

Table 1: Elasticities of relative structure type shares to housing density

	IHS relative medium size share	IHS relative large share
Log housing density	0.189***	0.092**
	(0.035)	(0.033)
Num. obs.	56 450	56 450
R^2	0.349	0.197
Mean dep. var.	0.516	0.207
Std. errs.	CBSA	CBSA
City FE	X	X

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.01$

Note:
This table displays the results of regression equation (39) for medium size apartment buildings and large apartment buildings.

3.4 Model Fit

In this section I show how the model is able to match some key untargeted features of the data. I first focus on measures of regulation. It is difficult to measure regulation for several reasons. The first is that it is a multi-dimensional phenomenon. Most municipalities regulate land use, including housing supply, with zoning codes and building codes that can run into the hundreds of pages. These documents contain rules about what types of housing may legally be built in different areas, as well as very specific restrictions on

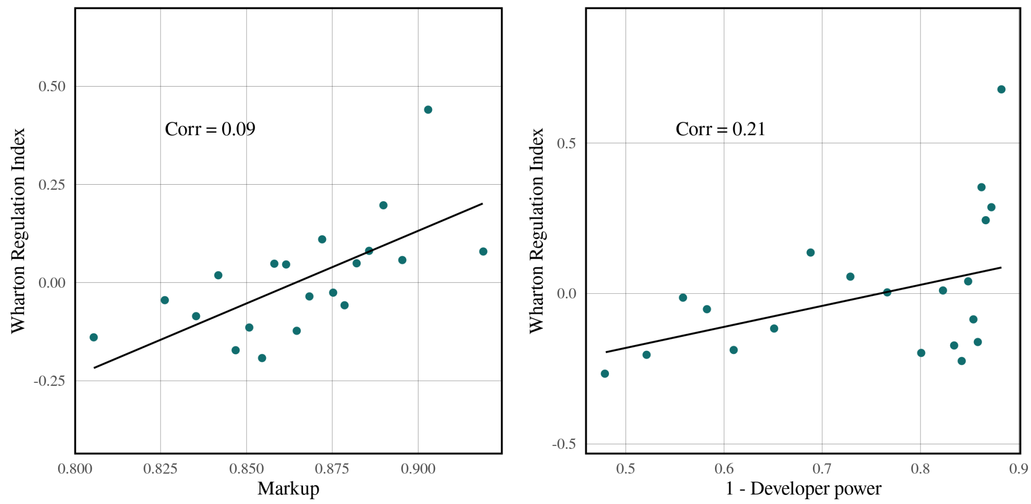

Figure 2: Binned scatter plots of model implied regulation measures and the Wharton Regulation Index

their physical form: minimum setbacks, minimum lot sizes, maximum floor area ratios, and many others. While the combined effect of these rules does control the amount of housing that gets built, it is difficult to summarise this in a single number.

Not only are these regulations very high dimensional, they are also enacted differently in different states. There is no federal database where all municipal zoning codes are indexed, and therefore researchers who wish to study them at large scale must either collect them by hand, or purchase them from from proprietary data aggregators. Even then, it is very difficult to cover even close to all of the urban municipalities in the US. These two issues have meant there have been very few attempts to measure the tightness of regulation at the municipal level. The highest profile of these is the Wharton Residential Land Use Regulation Index (henceforth Wharton Regulation Index) (Gyourko et al., 2021). It was constructed from a large mail survey where nearly 2,500 municipalities were asked in detail about different aspects of their zoning codes. The results of this survey were then summarised into a single index to capture the overall tightness of regulation.

My approach provides two key measures of the tightness of regulation. The first is the markup of floorspace prices above marginal construction costs. As mentioned above, this follows in the tradition of Glaeser and Gyourko (2002), although they conduct their analysis at the metropolitan area level. The other measure is one minus the inferred level of developer power, $1 - \xi_n$. This should be thought of as the institutional predisposition towards regulation, and the markup as the realised level of regulation. Given that the Wharton Regulation Index is defined at the municipality level, I take an average of my

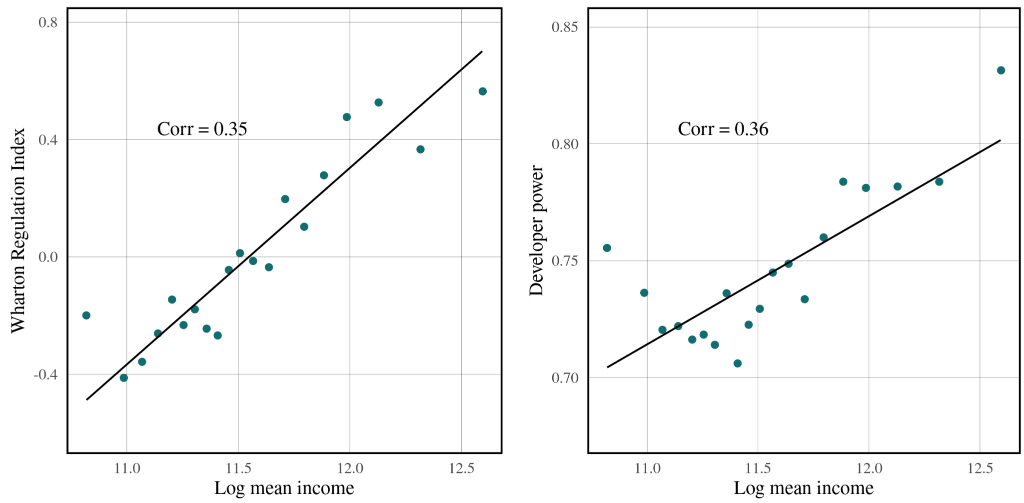

(a) Income and the Wharton Index

(b) Income and developer power

Figure 3: Binned scatter plots of regulation measures and income

census tract level measures, weighted by the number of households in each tract, in order to make them all comparable. Figure 2 shows binned scatter plots of both markups and developer power with the Wharton Regulation Index. In both cases, my model implied measures correlate positively with the observed measures. Municipalities with higher average markups and lower average developer power have higher levels regulation according to the Wharton Regulation Index. For reference, a recent paper by Bartik et al. (2024) uses large language models to analyse zoning texts, and its two main measures of regulation have correlations with the Wharton Index of 28% and 10% respectively.

A second key prediction of the model is the positive correlation between regulation and average neighbourhood income. Figure 3a establishes the fact empirically: richer municipalities have higher levels of the Wharton Regulation Index. Figure 3b then confirms this fact with the implied measure of developer power: neighbourhoods with lower developer power have higher average income. In the model, this relationship arises because areas with lower developer power have a higher propensity to regulate, regardless of the economic incentives to do so. They set tighter regulations, and therefore price out poorer households.

As mentioned above, two key advantages of my approach to measuring regulation are the extent of its coverage and its spatial resolution. The Wharton Regulation Index can only produce measures at the municipality level. However, there are other approaches to quantifying regulation at a finer spatial scale across many locations. Chief among these are the approaches of Song (2021) and Macek (2024), who use data on the size distribution of residential parcels and structural break detection algorithms to infer minimum lot sizes in different zoning districts within municipalities. This has the advantage of getting closer to measuring an objective feature of regulation that can be summarised easily; naturally, the downside is that it is unable to capture other dimensions of regulation. My approach takes the other side of the trade-off. I use the simple logic of the model to motivate markups and developer power as sufficient statistics for all of the complex ways that regulations and legal institutions work. This allows me to make claims about regulation at a fine spatial scale using accessible data, but has the clear downside of being further removed from direct data on regulation itself.

A helpful visualisation of the value of this approach is to map the implied measures of regulation. Figure 4 shows maps of developer power, ξ_n, for the three largest metropolitan areas in the US. The first clear pattern is that some municipal boundaries are visible. In particular, one can see that the central municipalities of each metropolitan area are inferred to have substantially higher developer power than the surrounding municipalities. At a stylised level, this is coming from two facts. The first is that these central municipalities are expensive to build in, since they are very dense, so they have relatively low markups. The second is that they have large stocks of valuable housing, which means new construction imposes high pecuniary externalities on the owners of existing housing. These municipalities therefore have strong incentives to block new development; the fact markups are nonetheless relatively low implies that developer power must be high.

The converse observation holds for the less dense, richer municipalities visible in the maps, such as Beverly Hills in Los Angeles or Evanston in Chicago. Their low density means that the marginal cost of building a home there is very low, and yet prices are very high. Developers therefore have a high willingness to pay for permits to build more housing there. Even the financial interests of the property owners are not enough to rationalise this without low political power of the developers. In general, with the exception of the five boroughs of New York City, one can see that developer power is almost everywhere inferred to be less than one half; that is, property owners are inferred to have most of the political power in the urban US. This is an insight that fits folk wisdom about the pervasiveness of 'NIMBYism' in municipal decision-making.

4 Counterfactuals

I now use the quantified model to study how the rise in the college wage premium, and subsequent adjustment of regulation, has increased income segregation. In section 4.1, I provide an overview of the solution procedures that I use to solve the model and isolate

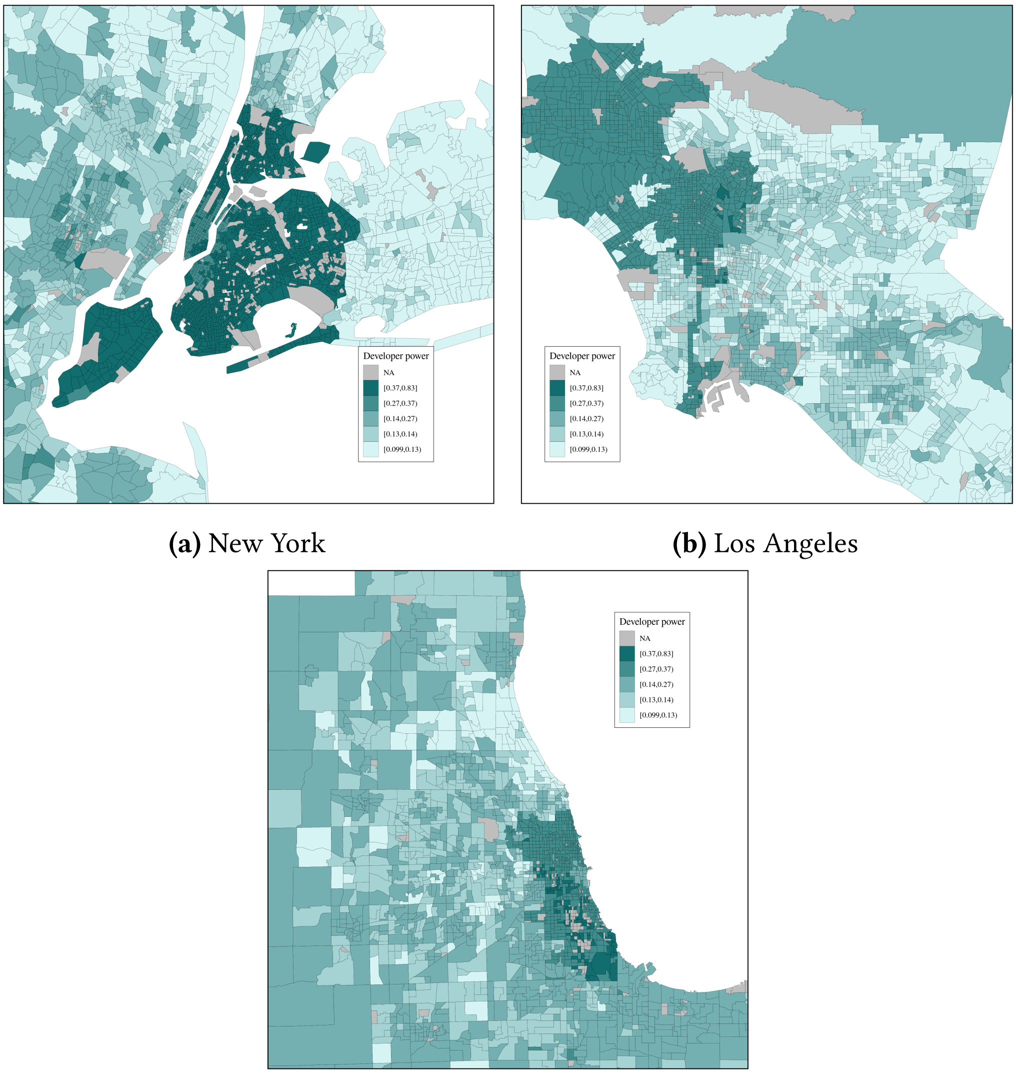

(a) New York

(b) Los Angeles

(c) Chicago

Figure 4: Maps of developer power, ξ_n

the effect of endogenous regulation. In section 4.2 I then describe the counterfactual exercises that I do, and their results.

4.1 Solving the Model

Here I provide an overview of the procedure that I use to solve the model. Importantly, recall from definition 1 that an equilibrium is nothing more than the labour markets and housing markets clearing in each neighbourhood, and therefore can be summarised by arrays of wages $\boldsymbol{w}$ and prices $\boldsymbol{p}$.

1. Start with an initial guess of wages $\boldsymbol{w}^0$ and prices $\boldsymbol{p}^0$. Set a convergence speed

parameter Ω.

2. Compute labour supply with equation (12), and the subsequent inverse labour demand of firms (i.e. their willingness to pay for this amount of labour) from equation (18). Call this inverse labour demand $\hat{\boldsymbol{w}}^0$.

3. Compute the first order condition of each municipality for each of their constituent neighbourhoods with equation (1), using the procedure described in section 3.2.5. Call the result of this $\boldsymbol{D}^0$.

4. Compute residual $R^0 \equiv \frac{1}{\sqrt{|\mathcal{N}|(1+|\mathcal{S}|)}} \left(\left|\left| \frac{\hat{\boldsymbol{w}}^0}{\boldsymbol{w}^0} \right|\right| + ||\boldsymbol{D}^0|| \right)$.

5. Update wages according to $w_n^1 \equiv w_n^0 \cdot \left(\frac{\hat{w}_n^0}{w_n^0} \right)^{\Omega}$.

6. Update prices according to $p_n^1 \equiv p_n^0 \cdot \exp\left(-\Omega \cdot D_n^0\right)$.

7. Iterate until R^i is less than some pre-established tolerance.

In section 2.7, I outline a simple version of my model, and show how the effect of a shock to the college wage premium can be decomposed into a 'direct effect' and an 'amplification effect' due to the endogenous adjustment of regulation. The direct effect is the result of the municipality making its decisions as if the income distribution in the neighbourhood were fixed. The amplification effect instead captures the effects of the municipality endogenously adjusting regulation to account for the changing income composition.

This suggests a quantitative analogue of the simple decomposition of these two effects. Given a shock to fundamentals, I first solve the full quantitative model using the method described above. I call this the 'full' solution. Any differences in outcomes between this solution and the baseline equilibrium can be thought of as the sum of the direct effect and the amplification effect.

I then solve the model again, this time holding fixed the income distribution in each neighbourhood that each municipality uses when it calculates its first order condition (equation 1). I call this the 'restricted' solution. Of course, the municipality is incorrect in its belief; as it changes the quantity of housing that it permits in each neighbourhood, households of different incomes move in and out, and change the local income distribution in doing so. I impose that the municipalities do not internalise this in their decision-making, and solve the rest of the model as usual. Any differences in outcomes between this restricted solution and the baseline equilibrium can therefore be thought of as only the direct effect. As a result, one can read the amplification effect from differences in outcomes between the full solution and the restricted solution.

4.2 Increasing the College Wage Premium

I conduct counterfactual exercises to explore how the rise in the college wage premium since 1980 has increased income segregation. I run these counterfactual exercises separately for the three largest metropolitan areas in the US: New York, Los Angeles, and Chicago. Together, these cities represent one eighth of the country's population, and one fifth of its GDP. Moreover, precisely because they are so large, they occupy an outsized position in popular discourse about income segregation. I therefore consider them an important benchmark; nonetheless, an advantage of this framework and its light data requirements is that such counterfactuals can be extended to nearly all of the metropolitan areas in the US without issue.[11] As such, I plan to do this in future work. Note that the model does not explicitly contain multiple cities; rather, it contains one city and an outside option representing the rest of the United States.

For each city, I conduct two solutions of the model. The first simply aims to 'rewind' the college wage premium back to 1980, to have an equilibrium to compare to. I am not able to observe this equilibrium directly, because the datasets needed to invert unobserved location fundamentals from the model are not available for the year 1980. Therefore, I hold everything about the 2017 equilibrium fixed, except I reduce the wages of high-skill households. In particular, I follow Cline and Kaymak (2025) who document a 75% rise in the college wage premium between 1980 and 2017. As such, I take the college wage premium in each neighbourhood that I recover from the 2017 equilibrium, and multiply it by $1/1.75$. I keep the wages of the low-skill households the same everywhere, so that this only represents a reduction in high-skill wages. This leaves me with a new array of wages, $\boldsymbol{w}^{1980}$. I then solve the model, imposing that wages must be equal to $\boldsymbol{w}^{1980}$. Implicitly, this means choosing an array of productivities $\boldsymbol{a}^{1980}$ so that they exactly rationalise $\boldsymbol{w}^{1980}$ as part of the equilibrium of the model.

I then perform a restricted solution of the model as described in section 4.1. In this case, I impose that municipalities make their regulatory decisions based on the local income distributions recovered in the previous solution. However, productivities are returned to their 2017 levels, which means wages for high-skill households rise again. This restricted solution should be thought of as the direct effect of the increased college

[11]The main data constraint is the RSMeans construction cost data, which does not provide location adjustment factors for some smaller cities.

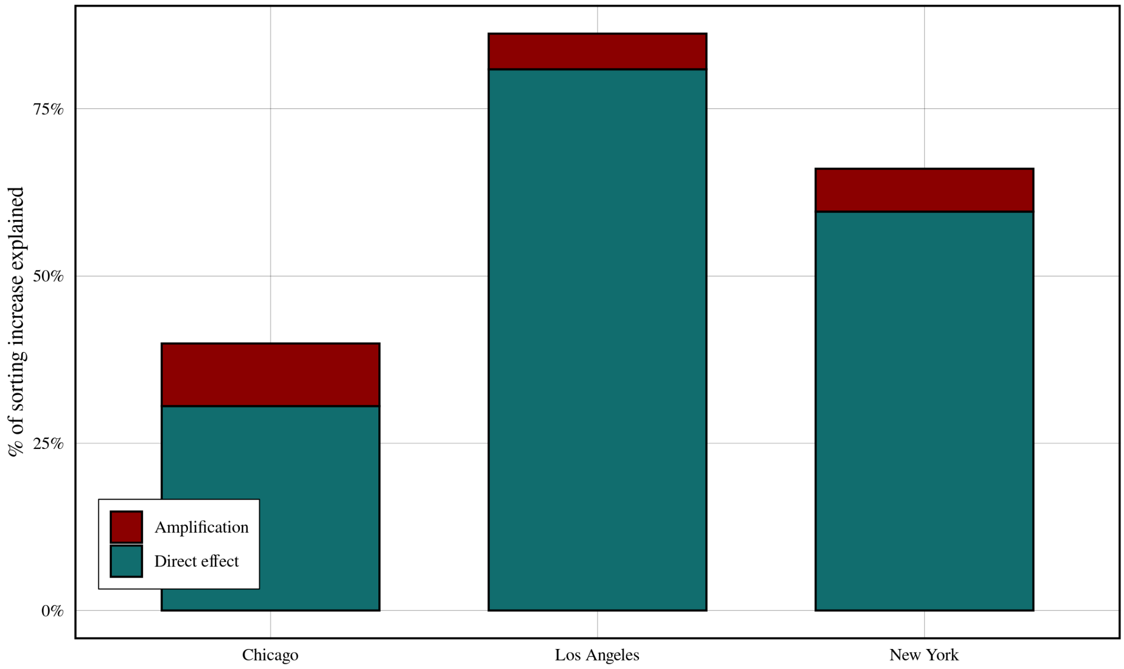

Figure 5: Observed segregation increases explained by the increased college wage premium

wage premium: high-skill households become richer and change their sorting patterns as a result, but municipalities do not change the stringency of their regulation to reflect this.

To measure the outcomes of these exercises, I use a standard measure of income sorting known as the Theil index. The index is defined as follows:

$$H \equiv \frac{\sum_{n \in \mathcal{N}} \sum_{w \in \mathcal{W}} \mathbb{P}(n, w) \ln \left(\frac{\mathbb{P}(n,w)}{\mathbb{P}(n)\mathbb{P}(w)} \right)}{\sum_{w \in \mathcal{W}} \mathbb{P}(w) \ln \left(\frac{1}{\mathbb{P}(w)} \right)} \in [0, 1] \tag{40}$$

where $n \in \mathcal{N}$ are neighbourhoods and $w \in \mathcal{W}$ are income quantiles. Formally, it is the normalised Kullback-Leibler divergence of the joint distribution of income and location choices from the product of their marginal distributions. This quantifies the informational content of location choices about income, and vice versa. If a city is very segregated by income, knowing the neighbourhood that a household lives in provides a great deal of information about its income.

Figure 5 shows the effect of the wage premium shock on sorting in the three metropolitan areas that I focus on. Depending on the city, the effects of the shock explain 40 to 86% of the observed increase in income segregation. The primary reason that the shock is less able to explain the increase in segregation in Chicago is that Chicago underwent a much higher increase in segregation during this time period than the other two cities.

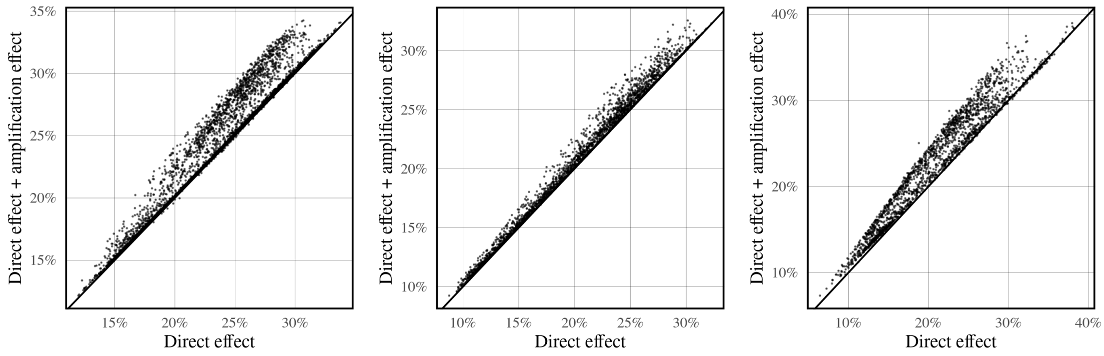

(a) New York **(b)** Los Angeles **(c)** Chicago

Figure 6: Income growth amplification with skill biased technical change

The predicted effect is remarkably similar among the three cities, suggesting that this may be due to an idiosyncratic set of other shocks that only affected Chicago.

Consistent with the predictions of the model, regulation amplifies income segregation. Between 6 and 29% of the total effect in Figure 5 comes from the endogenous adjustment of regulation. This also holds when one looks at a finer spatial scale. Figure 6 plots the growth in neighbourhood average income at the neighbourhood level with the direct effect on the horizontal axis and the sum of the direct and amplification effects on the vertical axis. The first pattern to note is that the points nearly all lie above the 45 degree line; that is, the amplification effect is positive, and often quite substantial. The second pattern is that the points are not uniformly distributed. There are clusters with higher and lower amounts of amplification, which correspond to municipalities with differing levels of developer power. This highlights the fact that the overall amplification of income sorting in Figure 5 masks substantial heterogeneity in the response of regulation to the wage premium shock.

The binned scatter plots in Figure 7 explore this heterogeneity further. They plot the size of the amplification effect on income growth, relative to the direct effect, as a function of developer power. The bulk of the amplification effect is coming from the neighbourhoods with the lowest developer power, and it decays quite rapidly as developer power increases. This should be intuitive, given the structure of the model. Neighbourhoods where the interests of the owners of inframarginal housing are more powerful will be regulated more according to their wishes, rather than those of the marginal developer. As such, any changes in the price elasticity of housing demand will pass through more into changes in markups. This can be seen in equation (2); the relative political power of inframarginal property owners, $\frac{1-\xi_n}{\xi_n}$, acts as a multiplier in the relationship between externalities on property owners and markups.

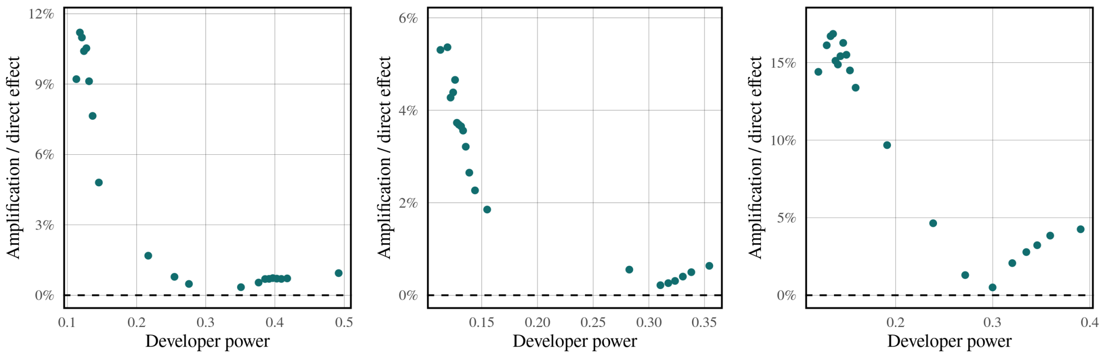

(a) New York **(b)** Los Angeles **(c)** Chicago

Figure 7: Developer power and the size of the amplification effect

5 Conclusion

This paper studies how income inequality increases income segregation, and how housing supply regulation endogenously adjusts to amplify this. I propose a simple theory of endogenous housing supply regulation, and embed it in a quantitative urban model of income sorting. I capture the fundamental tension between owners of existing housing and developers of new housing, and show how the resulting regulation exacerbates differences in average income across space. As neighbourhoods experience income shocks, the elasticity of demand for housing shifts, and this changes the marginal gains that property owners receive from restricting new construction. As a result, municipalities change regulation and cause further changes in average income. The result is a city that is more spatially unequal.

I show how this model can be quantified in a computationally tractable way, and bring it to bear with publicly available data in the US. I recover a novel measure of the implied weights on the interests of inframarginal property owners and marginal developers at the census tract level for nearly the entire urban US. These weights correlate well with the canonical existing measure of regulation, the Wharton Regulation Index, despite not using any data on regulation in computing them. Using these recovered weights and the rest of the quantified model, I show how the rise in the college wage premium since 1980 can explain 40-86% of the observed increase in income segregation, and that 6-29% of this effect is the result of the endogenous adjustment of regulation.

The framework that I outline in this paper suggests different lines of future research. For one, I believe that this modelling framework can be explored further for policy implications. Zoning reform and housing affordability are currently issues that are at the forefront of public discourse more so than they have been in the past, and yet they are not easy to make progress on. My model might be able both to help understand the politics of why there is so much localised resistance to zoning reform, and how states or the federal government might want to change local political institutions to better meet the housing demands of the coming decades.

Another clear avenue of research concerns ξ_n, developer power. I currently treat developer power as an exogenous fundamental of each neighbourhood. Of course, it is natural to ask whether it is truly exogenous, or if it too is an equilibrium object. It is certainly not obvious that physical geography should play any role in determining the relative political power of different groups.

Finally, in future I plant to incorporate tenure choice and dynamic wealth accumulation into the households' block of the model. The current version is what I consider to be the most parsimonious account of sorting and regulation that still captures what I consider to be the relevant features of the world. Moreover, this buys computational tractability which can then be used elsewhere in computing municipalities' optimality conditions for metropolitan areas with thousands of neighbourhoods. This is currently challenging to combine with a dynamic model, but I believe there is fruitful work to be done in that direction nonetheless.

References

Ahlfeldt, Gabriel M. et al. (2015). "The Economics of Density: Evidence From the Berlin Wall". In: ***Econometrica*** 83.6, pp. 2127–2189. ISSN: 1468-0262. DOI: 10.3982/ECTA10876.

Almagro, Milena and Tomás Domínguez-Iino (2025). "Location Sorting and Endogenous Amenities: Evidence From Amsterdam". In: ***Econometrica*** 93.3, pp. 1031–1071. ISSN: 1468-0262. DOI: 10.3982/ECTA21394.

Almagro, Milena et al. (2024). "Neighborhood Revitalization and Inequality: Evidence from Chicago's Public Housing Demolitions".

Babalievsky, Fil et al. (2024). "The Distributional Consequences of Residential Land Use Regulations".

Bagagli, Sara (2025). "The (Express)Way to Segregation: Evidence from Chicago".

Bartik, Alexander et al. (2024). "The Costs of Housing Regulation: Evidence from Generative Regulatory Measurement".

Bordeu, Olivia (2023). "Commuting Infrastructure in Fragmented Cities".

Brinkman, Jeffrey and Jeffrey Lin (Sept. 2022). "Freeway Revolts! The Quality of Life Effects of Highways". In: ***The Review of Economics and Statistics***, pp. 1–45. ISSN: 0034-6535. DOI: 10.1162/rest_a_01244.

Brooks, Leah and Byron Lutz (July 2019). "Vestiges of Transit: Urban Persistence at a Microscale". In: ***The Review of Economics and Statistics*** 101.3, pp. 385–399. ISSN: 0034-6535. DOI: 10.1162/rest_a_00817.

Cline, Alexander and Barış Kaymak (Mar. 2025). "Demand for College Labor in the 21st Century". In: ***Economic Commentary*** 2025-04. DOI: 10.26509/frbc-ec-202504.

Comin, Diego et al. (2021). "Structural Change With Long-Run Income and Price Effects". In: ***Econometrica*** 89.1, pp. 311–374. ISSN: 1468-0262. DOI: 10.3982/ECTA16317.

Couture, Victor and Jessie Handbury (Sept. 2020). "Urban Revival in America". In: ***Journal of Urban Economics*** 119, p. 103267. ISSN: 0094-1190. DOI: 10.1016/j.jue.2020.103267.

Couture, Victor et al. (Mar. 2024). "Income Growth and the Distributional Effects of Urban Spatial Sorting". In: ***The Review of Economic Studies*** 91.2, pp. 858–898. ISSN: 0034-6527. DOI: 10.1093/restud/rdad048.

D'Amico, Leonardo et al. (Nov. 2024). ***Why Has Construction Productivity Stagnated? The Role of Land-Use Regulation***. Working Paper 33188. National Bureau of Economic Research. DOI: 10.3386/w33188.

Desmet, Klaus and Jordan Rappaport (Mar. 2017). "The Settlement of the United States, 1800–2000: The Long Transition towards Gibrat's Law". In: ***Journal of Urban Economics***. Urbanization in Developing Countries: Past and Present 98, pp. 50–68. ISSN: 0094-1190. DOI: 10.1016/j.jue.2015.03.004.

Desmet, Klaus et al. (June 2018). "The Geography of Development". In: ***Journal of Political Economy*** 126.3, pp. 903–983. ISSN: 0022-3808. DOI: 10.1086/697084.

Diamond, Rebecca (Mar. 2016). "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000". In: ***American Economic Review*** 106.3, pp. 479–524. ISSN: 0002-8282. DOI: 10.1257/aer.20131706.

Dobbels, Gregory and Suren Tavakalov (2023). "Not in My Back Yard: The Local Political Economy of Residential Land-Use Regulations". In.

Duranton, Gilles and Diego Puga (2023). "Urban Growth and Its Aggregate Implications". In: ***Econometrica*** 91.6, pp. 2219–2259. ISSN: 1468-0262. DOI: 10.3982/ECTA17936.

Favilukis, Jack and Jaehee Song (2025). "Why Zoning Is Too Restrictive".

Finlay, John and Trevor Williams (2022). "Housing Demand, Inequality, and Spatial Sorting".

Fischel, William A. (Feb. 2005). ***The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies***. Cambridge, MA: Harvard University Press. ISBN: 978-0-674-01595-1.

Fogli, Alessandra et al. (2025). "The End of the American Dream? Inequality and Segregation in US Cities".

Ganong, Peter and Daniel Shoag (Nov. 2017). "Why Has Regional Income Convergence in the U.S. Declined?" In: ***Journal of Urban Economics*** 102, pp. 76–90. ISSN: 0094-1190. DOI: 10.1016/j.jue.2017.07.002.

Giannone, Elisa (2022). "Skill-Biased Technical Change and Regional Convergence".

Glaeser, Edward L. and Joseph Gyourko (Mar. 2002). ***The Impact of Zoning on Housing Affordability***. Working Paper. DOI: 10.3386/w8835. National Bureau of Economic Research: 8835.

Glaeser, Edward L. et al. (2005). "Why Have Housing Prices Gone Up?" In: ***American Economic Review*** 95.2, pp. 329–333.

Glaeser, Edward L. et al. (2008). "Why Do the Poor Live in Cities? The Role of Public Transportation". In: ***Journal of Urban Economics*** 63.1, pp. 1–24.

Gyourko, Joseph et al. (Mar. 2008). "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index". In: ***Urban Studies*** 45.3, pp. 693–729. ISSN: 0042-0980. DOI: 10.1177/0042098007087341.

Gyourko, Joseph et al. (July 2021). "The Local Residential Land Use Regulatory Environment across U.S. Housing Markets: Evidence from a New Wharton Index". In: ***Journal of Urban Economics*** 124, p. 103337. ISSN: 0094-1190. DOI: 10.1016/j.jue.2021.103337.

Harari, Mariaflavia (2024). "Residential Patterns and Local Public Goods in Urban Brazil".

He, Dongyang (2024). "Distributional Effects of Residential Zoning Policies: Insights from the Greater Boston Area".

Heblich, Stephan et al. (Nov. 2020). "The Making of the Modern Metropolis: Evidence from London*". In: ***The Quarterly Journal of Economics*** 135.4, pp. 2059–2133. ISSN: 0033-5533. DOI: 10.1093/qje/qjaa014.

Hempel, Alexander (2023). "The Impact of Greenbelts on Housing Markets: Evidence from Toronto".

Herkenhoff, Kyle F. et al. (Jan. 2018). "Tarnishing the Golden and Empire States: Land-use Restrictions and the U.S. Economic Slowdown". In: ***Journal of Monetary Economics***. Carnegie-Rochester-NYU Conference on Public Policy Held at the Stern School of Business at New York University 93, pp. 89–109. ISSN: 0304-3932. DOI: 10.1016/j.jmoneco.2017.11.001.

Hilber, Christian A. L. and Frédéric Robert-Nicoud (May 2013). “On the Origins of Land Use Regulations: Theory and Evidence from US Metro Areas”. In: ***Journal of Urban Economics*** 75, pp. 29–43. ISSN: 0094-1190. DOI: 10.1016/j.jue.2012.10.002.

Hoelzlein, Matthias (2023). “Two-Sided Sorting and Spatial Inequality in Cities”.

Hsieh, Chang Tai and Enrico Moretti (2019). “Housing Constraints and Spatial Misallocation”. In: ***American Economic Journal: Macroeconomics*** 11.2, pp. 1–39. ISSN: 19457715. DOI: 10.1257/MAC.20170388.

Kulka, Amrita (2019). “Sorting into Neighbourhoods: The Role of Minimum Lot Sizes”.

Kulka, Amrita et al. (2022). “How to Increase Housing Affordability? Understanding Local Deterrents to Building Multi-Family Housing”.

Lee, Sanghoon and Jeffrey Lin (Jan. 2018). “Natural Amenities, Neighbourhood Dynamics, and Persistence in the Spatial Distribution of Income”. In: ***The Review of Economic Studies*** 85.1, pp. 663–694. ISSN: 0034-6527. DOI: 10.1093/restud/rdx018.

Ma, Lei (2025). “Build What and for Whom? The Distributional Effects of Housing Supply”.

Macek, James (2024). “Housing Regulation and Neighborhood Sorting Across the United States”.

Mei, Mike (2022). “House Size and Household Size: The Distributional Consequences of the Minimum Lot Size Regulation”.

Monte, Ferdinando et al. (Dec. 2018). “Commuting, Migration, and Local Employment Elasticities”. In: ***American Economic Review*** 108.12, pp. 3855–3890. ISSN: 0002-8282. DOI: 10.1257/aer.20151507.

Nagpal Geetika; Ghandi, Sahil (2023). “Scaling Heights: Affordability Implications of Zoning Deregulation in India”.

Ortalo-Magné, François and Andrea Prat (Feb. 2014). “On the Political Economy of Urban Growth: Homeownership versus Affordability”. In: ***American Economic Journal: Microeconomics*** 6.1, pp. 154–181. ISSN: 1945-7669. DOI: 10.1257/mic.6.1.154.

Ouasbaa, Ghizlen et al. (2025). “When Developers Hold Office: Shaping Housing Supply through Local Politics”. In: ***IEB (Institut d’Economia de Barcelona) – Working Papers***.

Owens III, Raymond et al. (May 2020). “Rethinking Detroit”. In: ***American Economic Journal: Economic Policy*** 12.2, pp. 258–305. ISSN: 1945-7731. DOI: 10.1257/pol.20180651.

Parkhomenko, Andrii (Nov. 2023). “Local Causes and Aggregate Implications of Land Use Regulation”. In: ***Journal of Urban Economics*** 138, p. 103605. ISSN: 0094-1190. DOI: 10.1016/j.jue.2023.103605.

Rossi-Hansberg, Esteban (2004). “Optimal Urban Land Use and Zoning”. In: ***Review of Economic Dynamics*** 7, pp. 69–106. DOI: 10.1016/S1094-2025(03)00056-5.

Severen, Christopher (Mar. 2018). “Commuting, Labor, and Housing Market Effects of Mass Transportation: Welfare and Identification”. In: ***SSRN***. DOI: 10.21799/FRBP.WP.2018.14.

Shanks, Brendan (2021). “Land Use Regulations and Housing Development”.

Song, Jaehee (2021). “The Effects of Residential Zoning in U.S. Housing Markets”.

Trounstine, Jessica (May 2020). “The Geography of Inequality: How Land Use Regulation Produces Segregation”. In: ***American Political Science Review*** 114.2, pp. 443–455. ISSN: 0003-0554, 1537-5943. DOI: 10.1017/S0003055419000844.

Tsivanidis, Nick (2023). “Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá’s TransMilenio”.

A Appendix - Theory

A.1 Non-Homothetic CES Derivations

In this section I derive various expressions from the household's optimisation problem over housing and tradable goods consumption. I suppress subscripts for brevity. Their non-homothetic consumption aggregator is given by the following:

$$C(c,h) = \left(c^{\frac{\sigma-1}{\sigma}} + C(c,h)^{\frac{\nu}{\sigma}} h^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} \tag{A.1}$$

To derive the demands and price index, solve the cost minimisation problem with the following Lagrangian:

$$\mathcal{L} = -ph - c + \mu\left(\left(c^{\frac{\sigma-1}{\sigma}} + C^{\frac{\nu}{\sigma}} h^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} - C\right) \tag{A.2}$$

The first order conditions are the following:

$$p = \mu C^{\frac{1+\nu}{\sigma}} h^{-\frac{1}{\sigma}} \tag{A.3}$$

$$1 = \mu C^{\frac{1}{\sigma}} c^{-\frac{1}{\sigma}} \tag{A.4}$$

Therefore,

$$h = \mu^{\sigma} C^{1+\nu} p^{-\sigma} \tag{A.5}$$

$$c = \mu^{\sigma} C \tag{A.6}$$

Imposing that the Lagrangian's constraint binds and substituting in these expressions yields

$$C = \left[\mu^{\sigma-1} C^{\frac{\sigma-1}{\sigma}} + C^{\frac{\nu}{\sigma}} \mu^{\sigma-1} C^{\frac{(1+\nu)(\sigma-1)}{\sigma}} p^{1-\sigma}\right]^{\frac{\sigma-1}{\sigma}} \tag{A.7}$$

$$\mu = \left(1 + C^{\nu} p^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \tag{A.8}$$

Substituting (A.5) and (A.6) into the budget constraint definition, we have

$$w = ph + c \tag{A.9}$$

$$= p^{1-\sigma}\mu^{\sigma} C^{1+\nu} + \mu^{\sigma} C \tag{A.10}$$

$$= \mu^{\sigma} C \left(p^{1-\sigma} C^{\nu} + 1\right) \tag{A.11}$$

$$= \mu^{\sigma}\mu^{1-\sigma} C = \mu C \tag{A.12}$$

$$C = \frac{w}{\mu} \tag{A.13}$$

The Lagrange multiplier therefore becomes the price index, i.e. the shadow savings of allowing a marginal unit less of utility.

$$\mu = P = \left(1 + \left(\frac{w}{P}\right)^{\nu} p^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \tag{A.14}$$

Housing demand for an individual household becomes

$$h(p,w) = \left(\frac{p}{P(p,w)}\right)^{-\sigma}\left(\frac{w}{P(p,w)}\right)^{1+\nu} \tag{A.15}$$

A.2 Proof of Lemma 1

The non-homothetic price index takes the following form, suppressing subscripts for brevity:

$$P(p,w) = \left(1 + \left(\frac{w}{P(p,w)}\right)^{\nu} p^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \tag{A.16}$$

Define

$$G(p,w) = \left(\frac{w}{P(p,w)}\right)^{\nu} p^{1-\sigma} \tag{A.17}$$

Taking the elasticity with respect to p,

$$\frac{\partial \ln P}{\partial \ln p} = \frac{1}{1-\sigma}\frac{G(p,w)}{1+G(p,w)}\left(-\nu\frac{\partial \ln P}{\partial \ln p} + 1 - \sigma\right) \tag{A.18}$$

Taking the elasticity again, this time with respect to w,

$$\begin{aligned}\frac{\partial^2 \ln P}{\partial \ln p \partial \ln w} &= \frac{1}{1-\sigma}\left[\frac{G(p,w)}{1+G(p,w)}\left(\frac{\partial \ln G}{\partial \ln w} - \frac{G(p,w)}{1+G(p,w)}\frac{\partial \ln G}{\partial \ln w}\right)\left(-\nu\frac{\partial \ln P}{\partial \ln p} + 1 - \sigma\right)\right. \\ &\quad \left. + \frac{G(p,w)}{1+G(p,w)}\left(-\nu\frac{\partial^2 \ln P}{\partial \ln p \partial \ln w}\right)\right] &\text{(A.19)} \\ &= \frac{1}{1-\sigma}\frac{G(p,w)}{1+G(p,w)}\left[\frac{\nu}{1+G(p,w)}\left(1 - \frac{\partial \ln P}{\partial \ln w}\right) - \nu\frac{\partial^2 \ln P}{\partial \ln p \partial \ln w}\right] &\text{(A.20)} \\ \frac{\partial^2 \ln P}{\partial \ln p \partial \ln w} &= \frac{Y(p,w)}{1+Y(p,w)}\frac{1}{1+G(p,w)}\left(1 - \frac{\partial \ln P}{\partial \ln w}\right) &\text{(A.21)}\end{aligned}$$

where

$$Y(p,w) = \frac{\nu}{1-\sigma}\frac{G(p,w)}{1+G(p,w)} \tag{A.22}$$

Taking instead the elasticity of P with respect to w,

$$\begin{aligned}\frac{\partial \ln P}{\partial \ln w} &= Y(p,w)\left(1 - \frac{\partial \ln P}{\partial \ln w}\right) &\text{(A.23)} \\ &= \frac{Y(p,w)}{1+Y(p,w)} &\text{(A.24)}\end{aligned}$$

Therefore,

$$\frac{\partial^2 \ln P}{\partial \ln p \partial \ln w} = \frac{Y(p,w)}{(1+Y(p,w))^2}\frac{1}{1+G(p,w)} \tag{A.25}$$

Since $G(p,w) > 0$ for all $p, w > 0$, this expression is strictly negative if and only if $Y(p,w) < 0$. $Y(p,w)$ in turn, is the product of two terms. The second, $\frac{G(p,w)}{1+G(p,w)}$ is always strictly positive by the same logic as before. If, as in the statement of the lemma, $\nu < 0$ and $\sigma \in (0, 1)$, then the first term is strictly negative, and $Y(p, w) < 0$. Note that these conditions are satisfied by the parameters used in the quantification. As such,

$$\frac{\partial^2 \ln P}{\partial \ln p \partial \ln w} < 0 \tag{A.26}$$

$\square$

B Appendix - Empirics

Parameter	Description	Value	Source
γ	Within-city preference dispersion	4	Midpoint of literature estimates
ζ	Cross-city preference dispersion	2	Monte et al. (2018)
κ	Utility semi-elasticity from commuting time	0.01	Ahlfeldt et al. (2015)
ψ_L	Spillovers on the low-skilled	0.1	Macek (2024)
ψ_H	Spillovers on the high-skilled	0.3	
σ	EoS btw. housing and tradable consumption	0.5	Finlay and Williams (2022)
ν	Income effect on housing share	-0.3	
ρ	EoS in production btw. low- and high-skill	1.66	Diamond (2016)
δ	Land share in production	0.9	Desmet and Rappaport (2017)
$\daleth_2$	Elasticity of mid-size apt. buildings w.r.t. density	0.19	Own estimate
$\daleth_3$	Elasticity of large apt. buildings w.r.t. density	0.09	

Table B.1: Summary of location-independent parameters

B.1 Summary Statistics

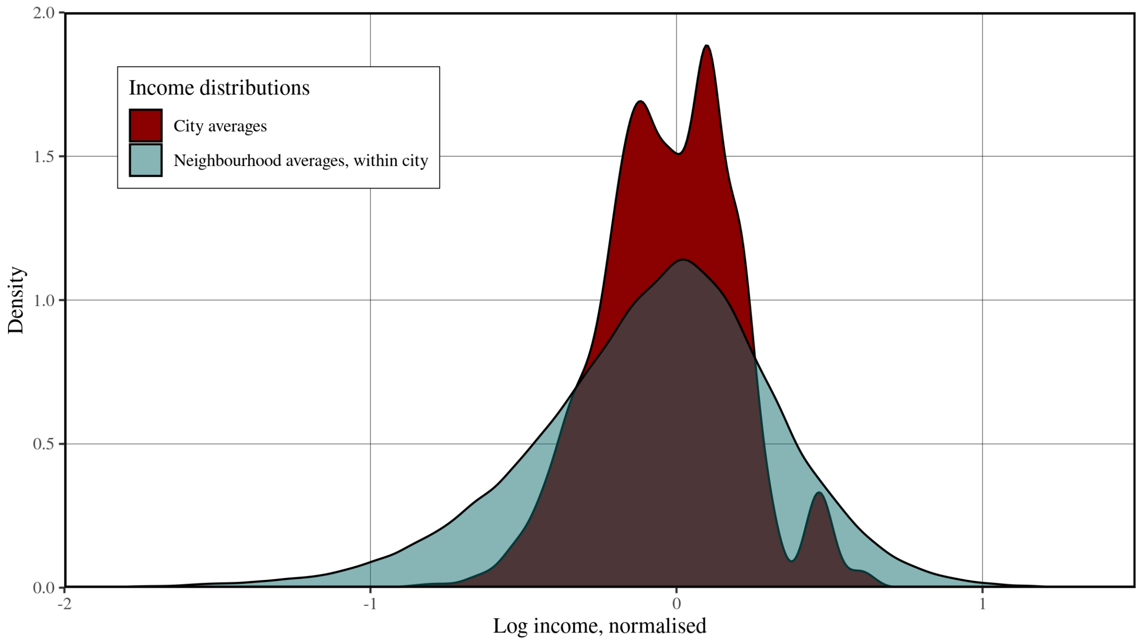

Figure B.1: Income distributions within and across cities.

Note: This figure shows the distribution of income across neighbourhoods within cities (in blue) and the distribution of average income across cities (in red).

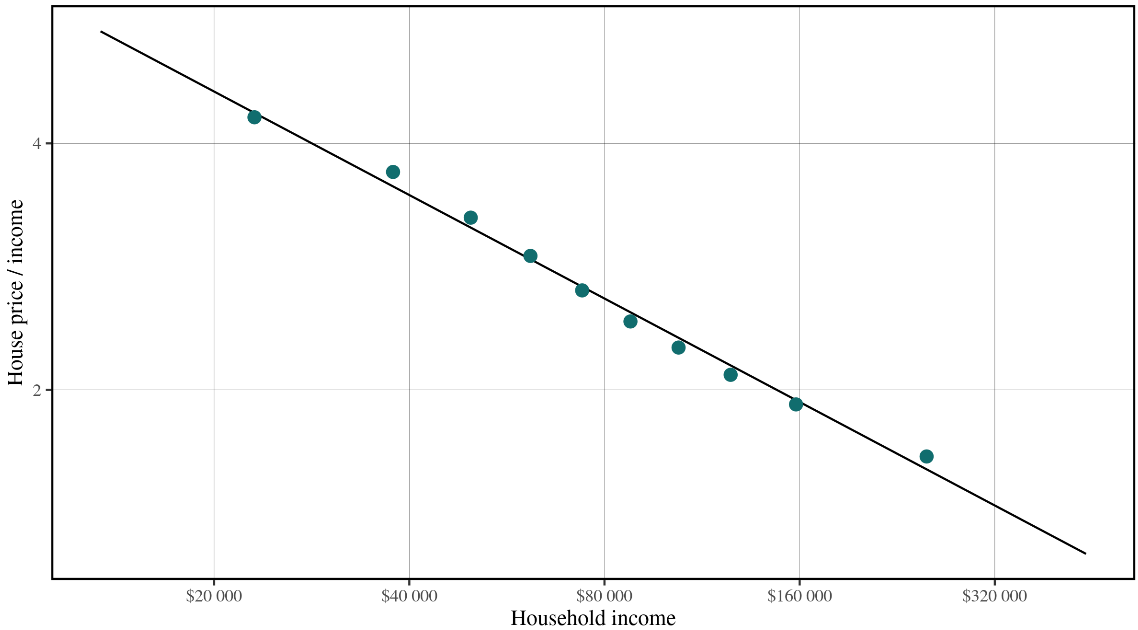

Figure B.2: Household income and housing prices

Note: This figure is a binned scatter plot of household income against the ratio of the household's house price to its income. It is computed from the 2017 ACS microdata sample, accessed via IPUMS.

B.2 Baseline Descriptives for Three Cities

Below I plot descriptive statistics for the three cities in the main part of my analysis.

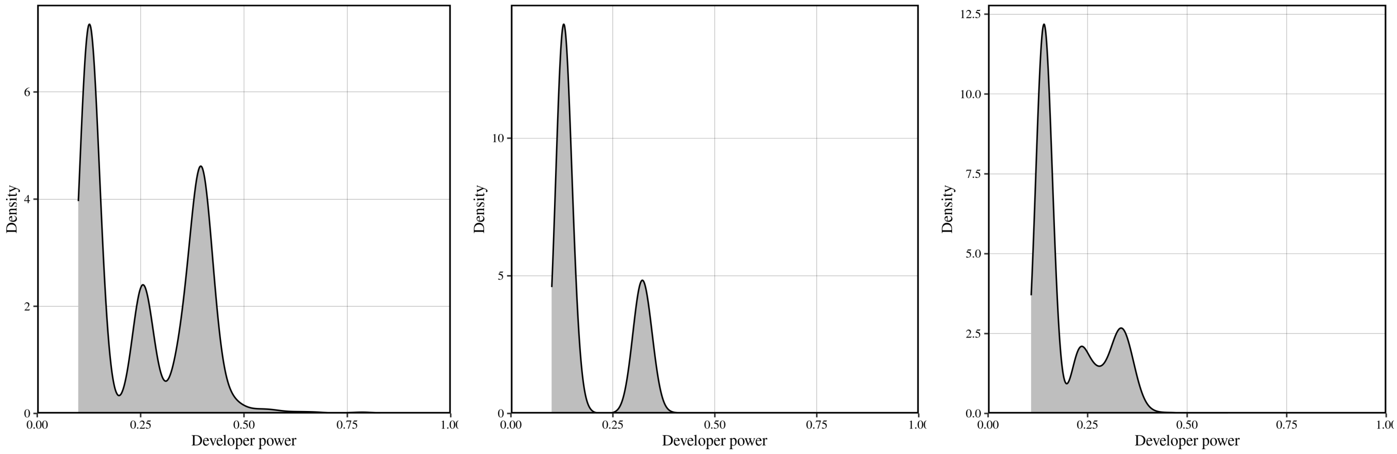

(a) New York City **(b)** Los Angeles **(c)** Chicago

Figure B.3: Developer power

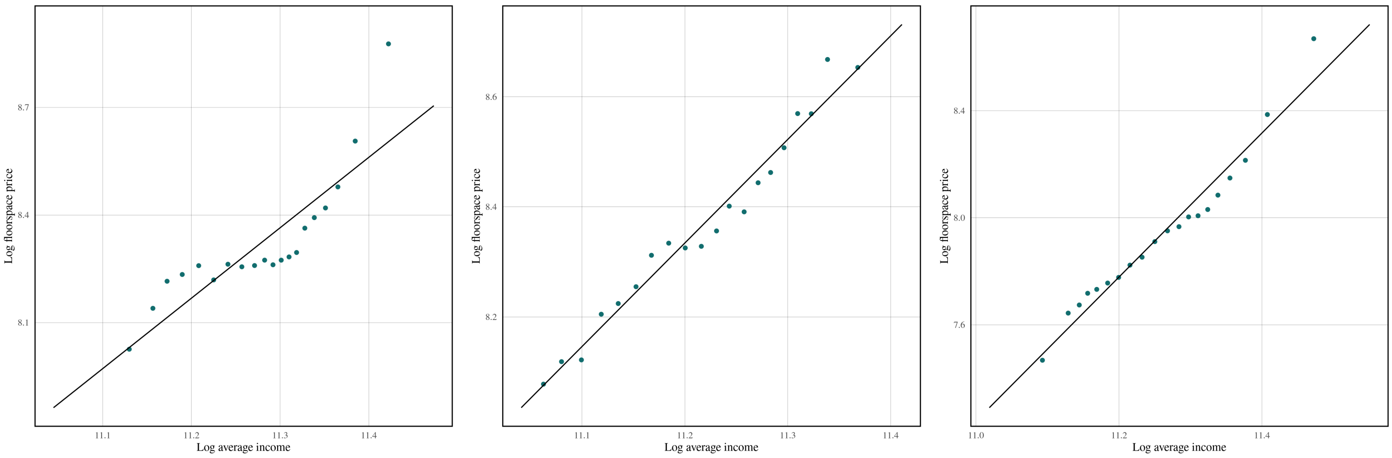

(a) New York City **(b)** Los Angeles **(c)** Chicago

Figure B.4: Neighbourhood income and prices

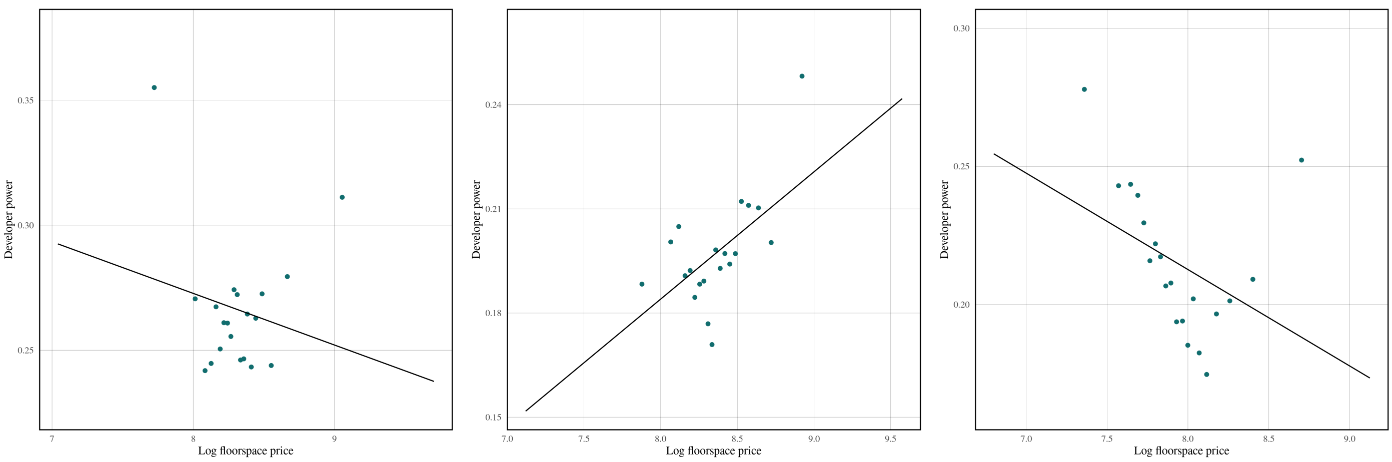

(a) New York City **(b)** Los Angeles **(c)** Chicago

Figure B.5: Prices and developer power

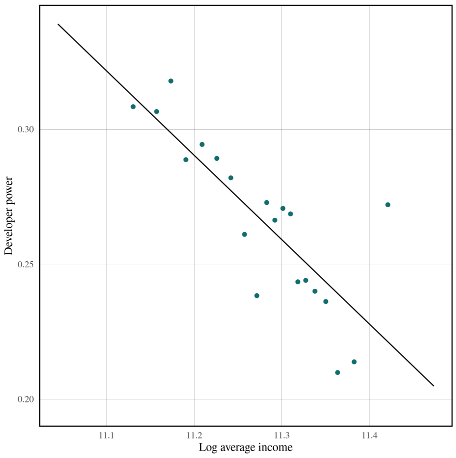

(a) New York City

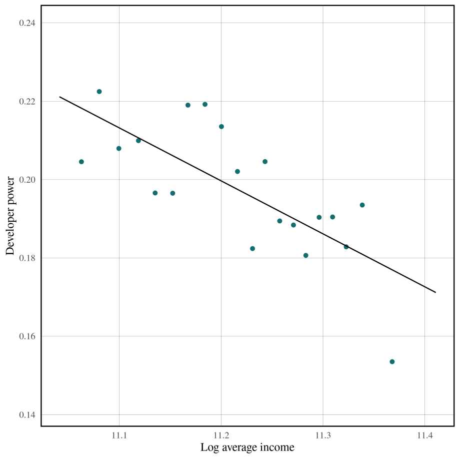

(b) Los Angeles

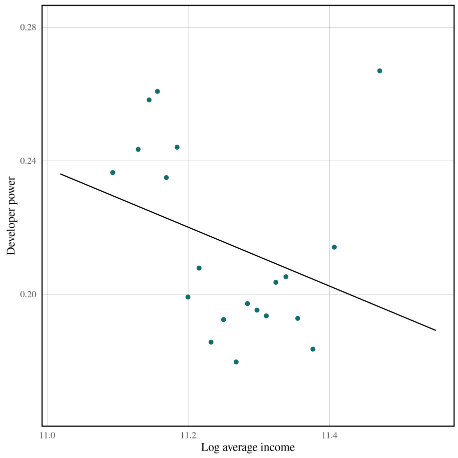

(c) Chicago

Figure B.6: Neighbourhood income and developer power

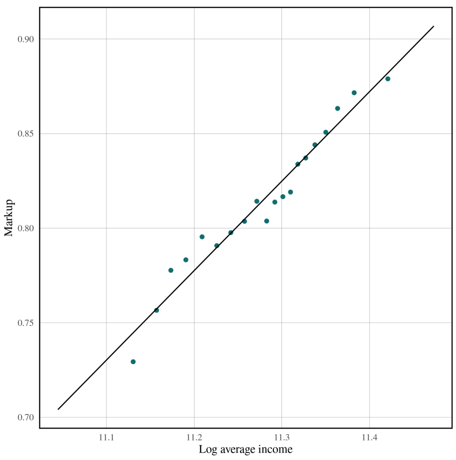

(a) New York City

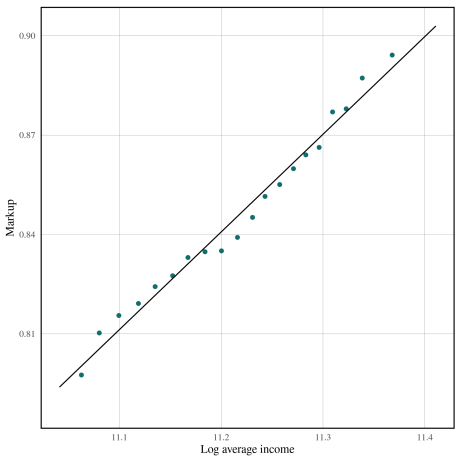

(b) Los Angeles

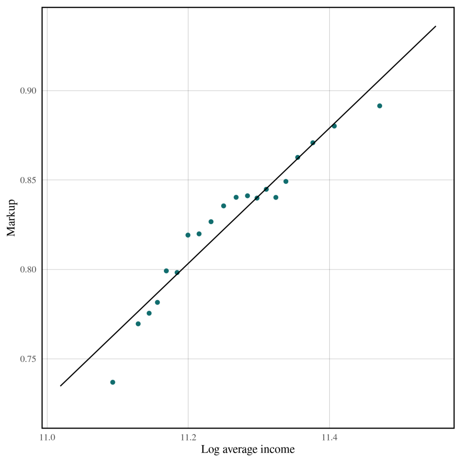

(c) Chicago

Figure B.7: Neighbourhood income and markups

	Developer power		
	(1)	(2)	(3)
Log mean income	-0.559^{***}	-0.472^{***}	-0.387^{***}
	(0.030)	(0.027)	(0.026)
Num. obs.	56 294	56 294	56 294
R^2	0.232	0.490	0.684
Mean dep. var.	0.702	0.702	0.702
Std. errs.	Municipality	Municipality	Municipality
City FE		X	
Muni. FE			X

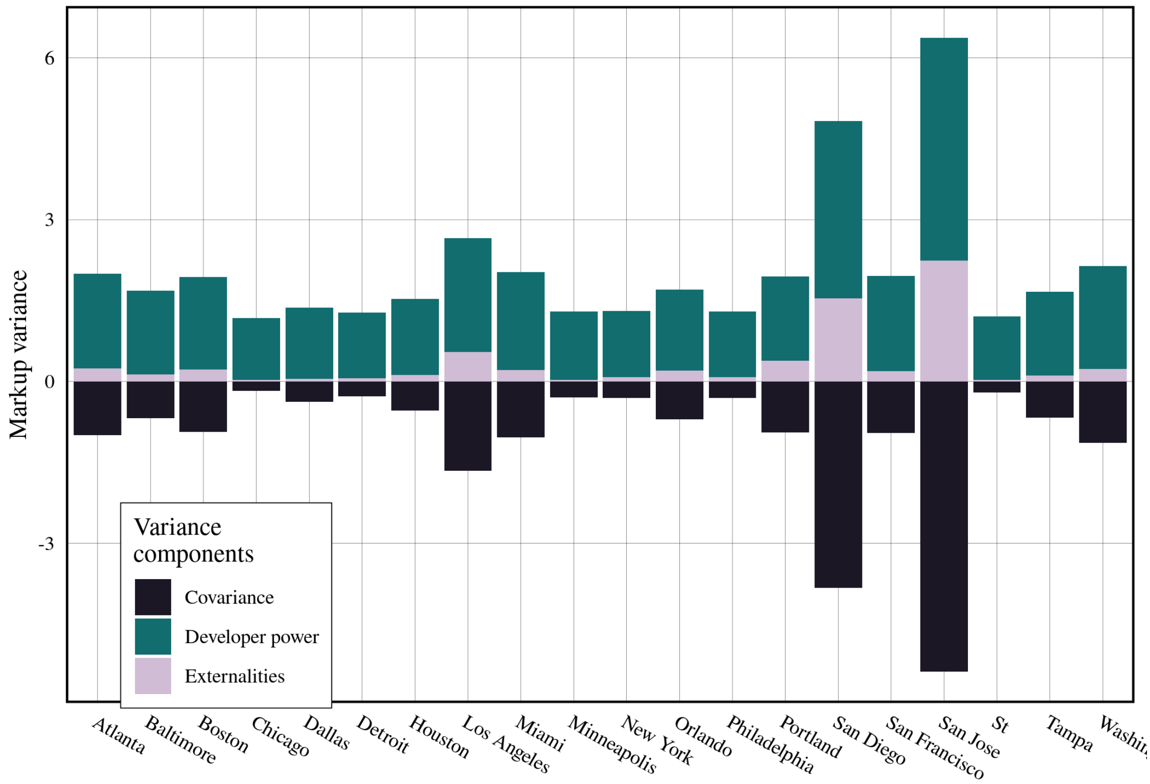

Figure B.8: Variance decomposition of $\ln((p_n - c'_n(h_n))/p_n)$ in the twenty largest metro areas

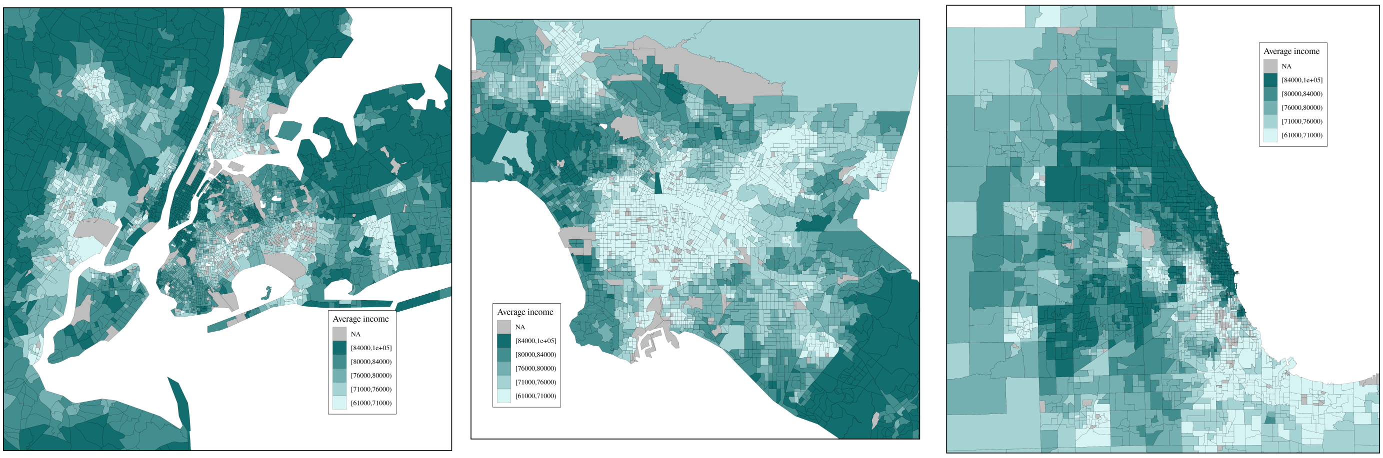

(a) New York City **(b)** Los Angeles **(c)** Chicago

Figure B.9: Average income

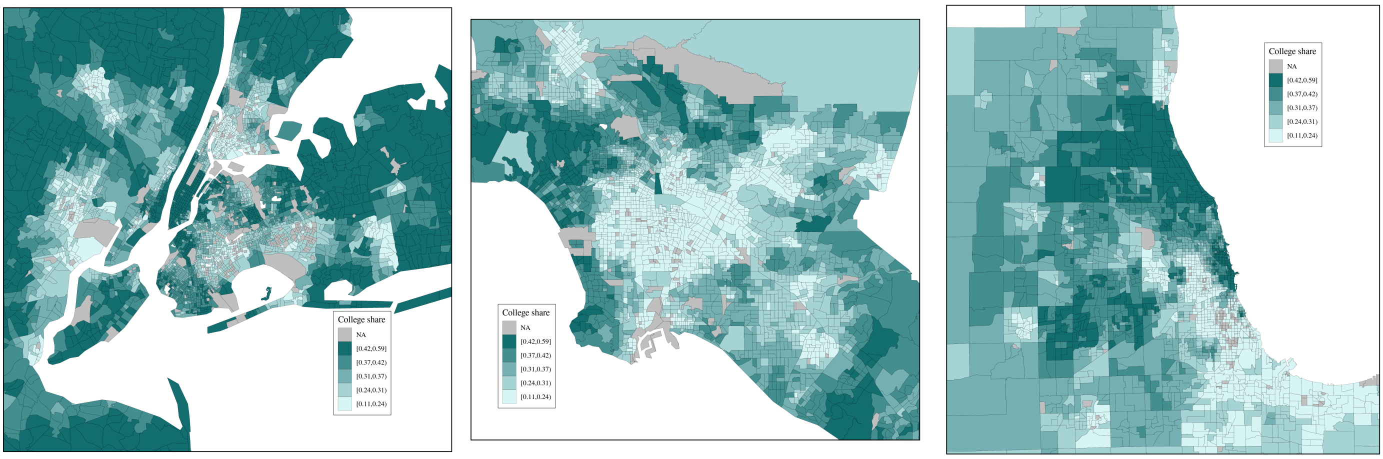

(a) New York City **(b)** Los Angeles **(c)** Chicago

Figure B.10: College share

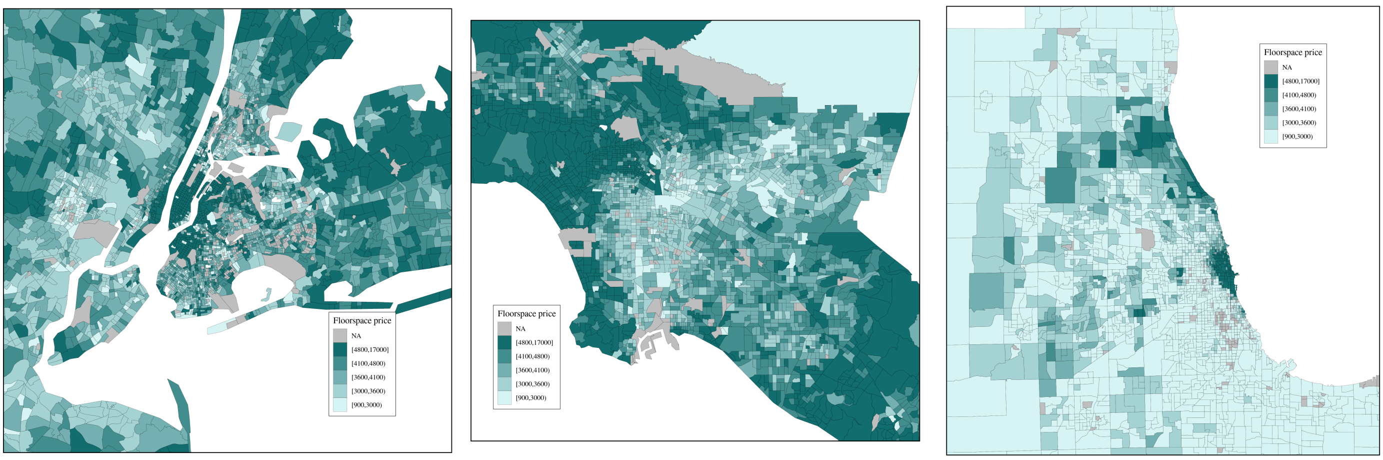

(a) New York City **(b)** Los Angeles **(c)** Chicago

Figure B.11: Floorspace price

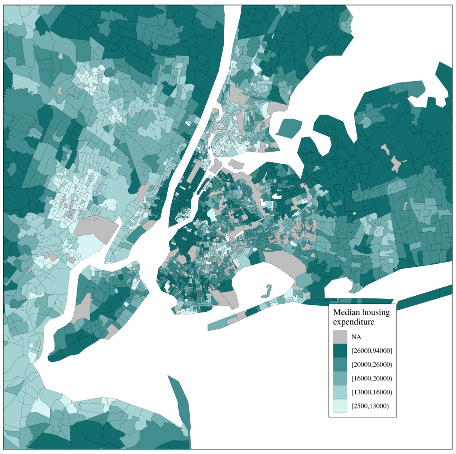

(a) New York City

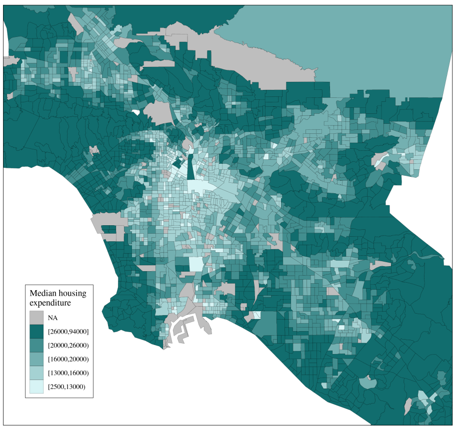

(b) Los Angeles

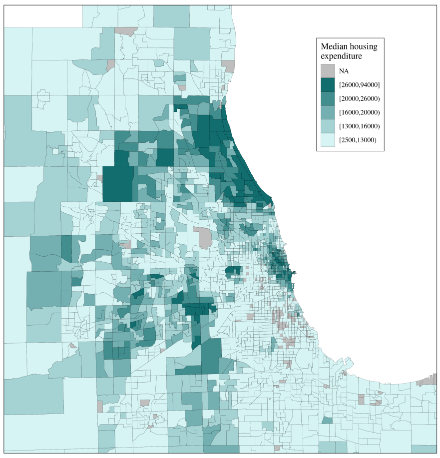

(c) Chicago

Figure B.12: Median housing expenditure

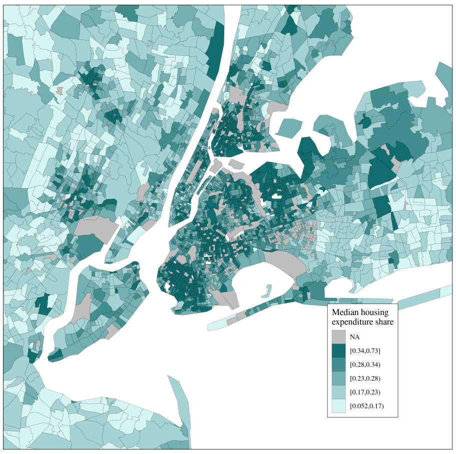

(a) New York City

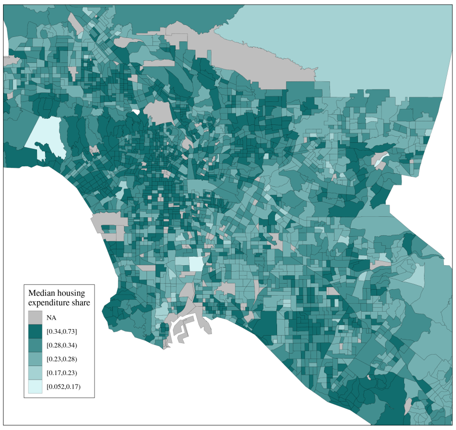

(b) Los Angeles

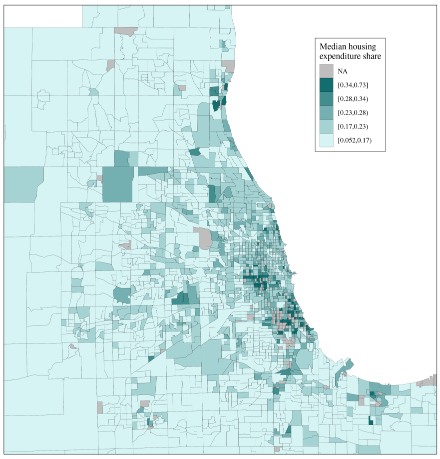

(c) Chicago

Figure B.13: Median housing expenditure share

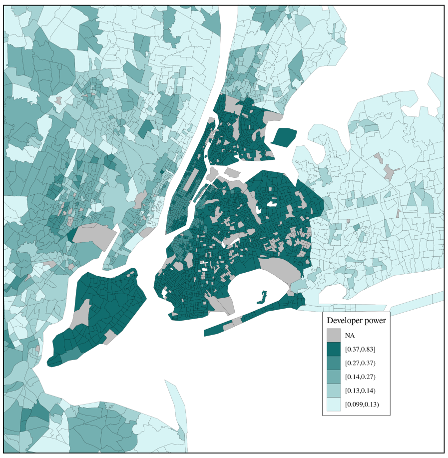

(a) New York City

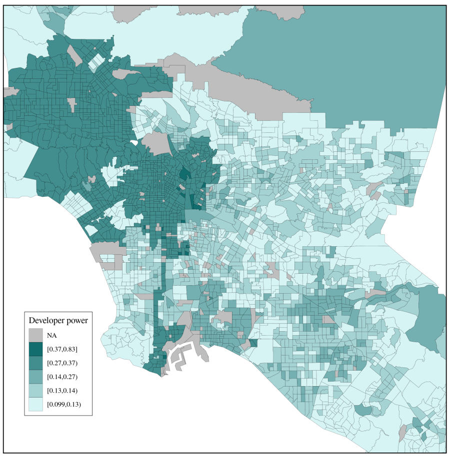

(b) Los Angeles

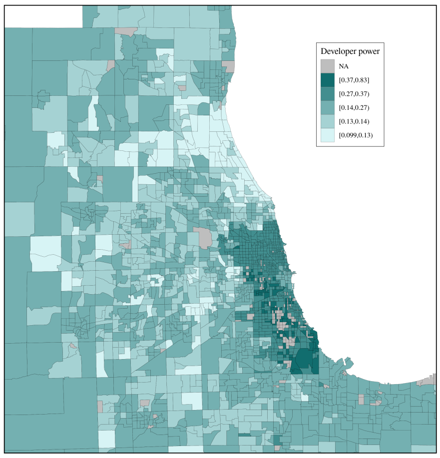

(c) Chicago

Figure B.14: Developer power

B.3 Increasing the College Wage Premium

B.3.1 Bar Plots of Citywide Aggregates

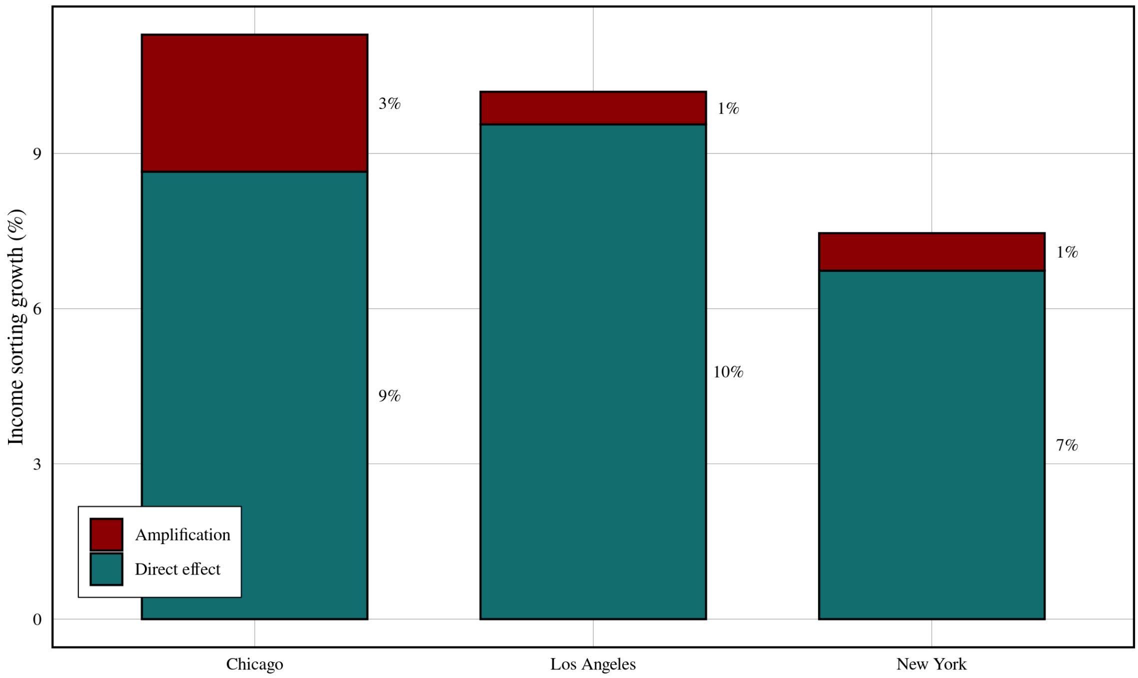

Figure B.15: Effect of skill biased technical change on sorting (Theil index)

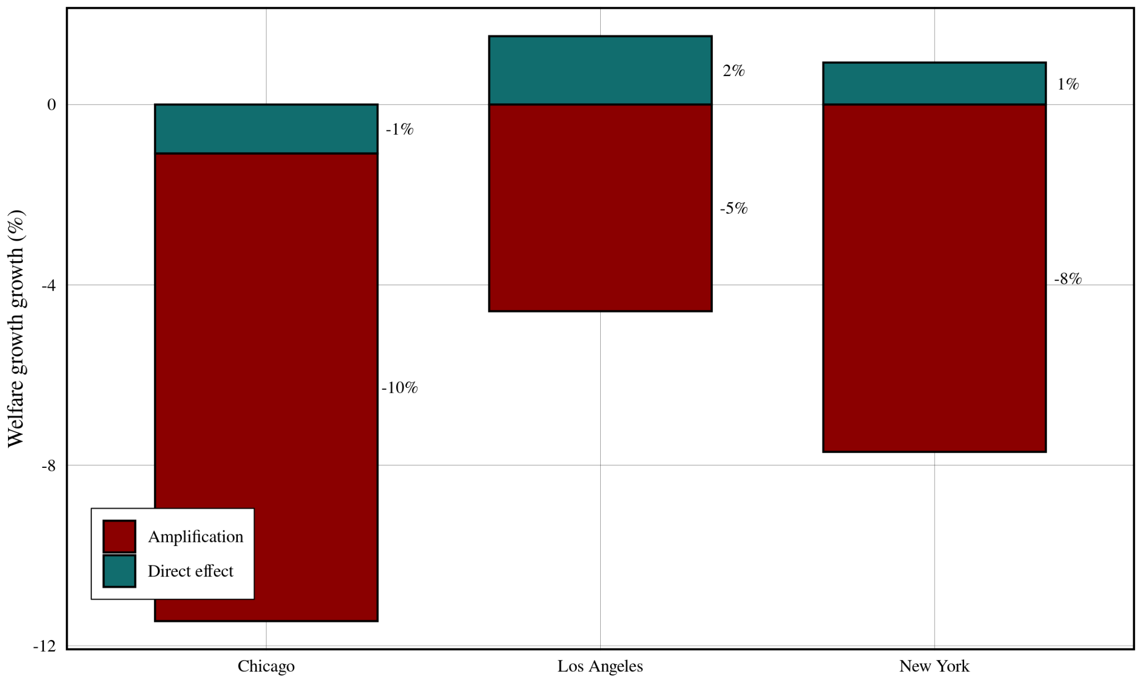

Figure B.16: Effect of skill biased technical change on low skill welfare

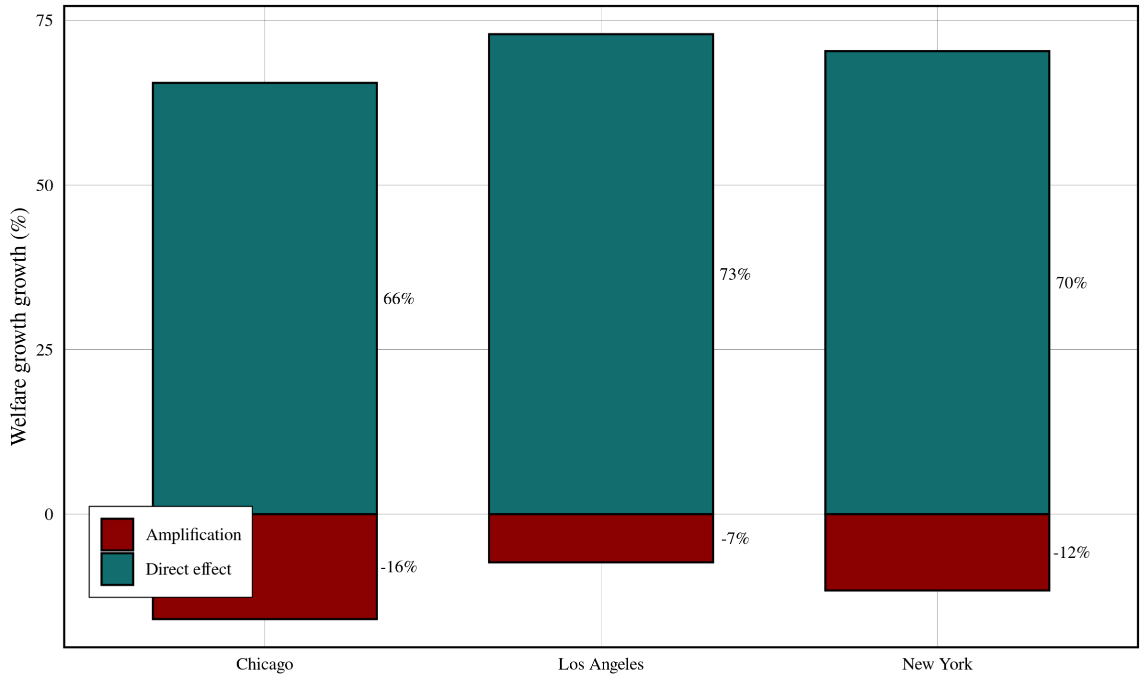

Figure B.17: Effect of skill biased technical change on high skill welfare

B.3.2 Scatter Plots (Changes)

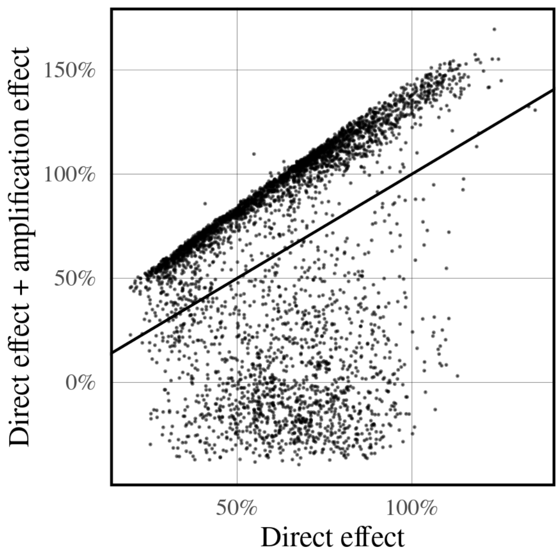

(a) New York

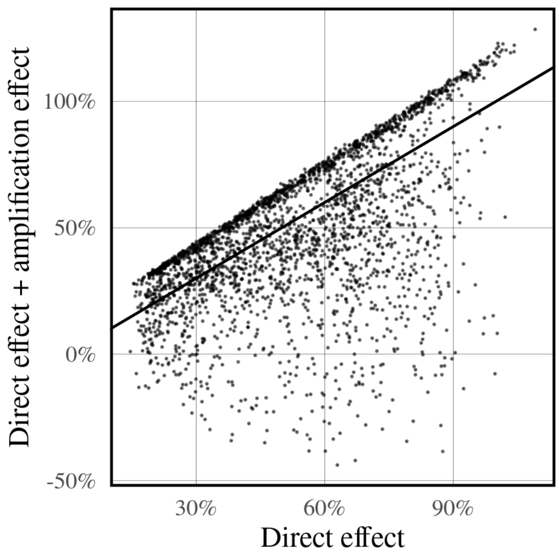

(b) Los Angeles

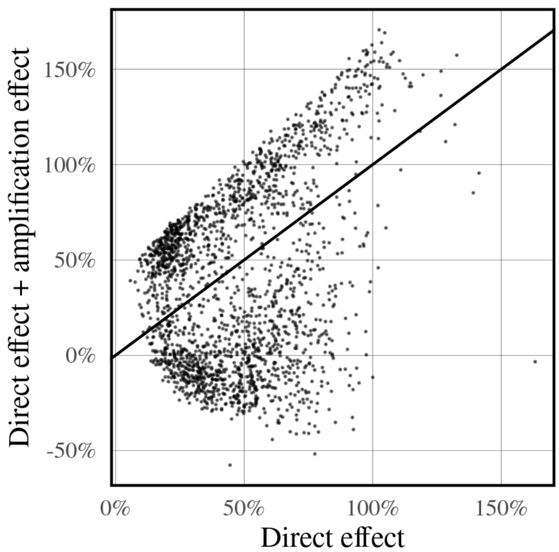

(c) Chicago

Figure B.18: Housing growth amplification with skill biased technical change

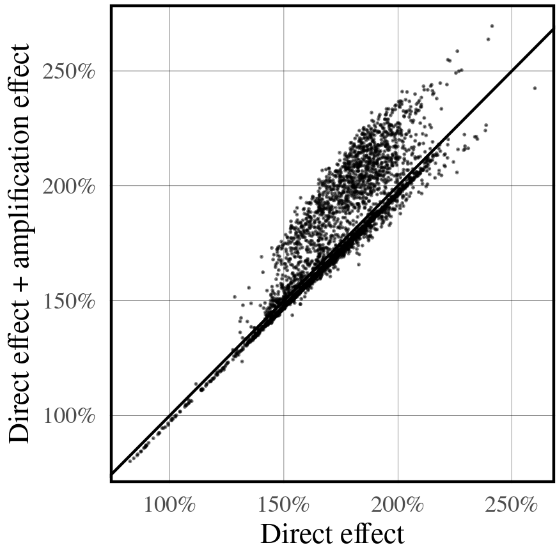

(a) New York

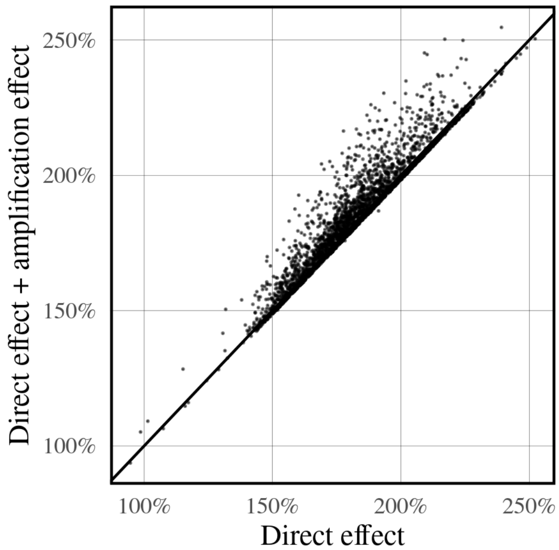

(b) Los Angeles

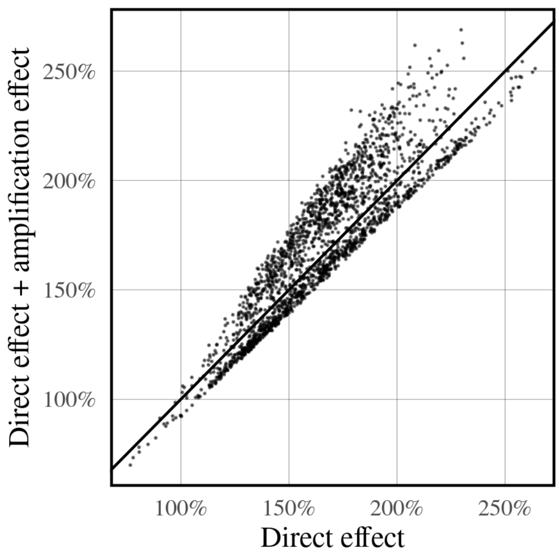

(c) Chicago

Figure B.19: Skill ratio growth amplification with skill biased technical change

B.4 Discussion of Alternate Approaches to Computing Cross-Price Elasticities

The simplest approach would be to use a finite-differences approximation. That is, for each pair of neighbourhoods (n', n), computing

$$\frac{dh_n}{dp_{n'}} \approx \frac{h_n\left(p_1, p_2, \ldots, p_{n'} + \Delta, \ldots, p_{|\mathcal{N}|}\right) - h_n\left(p_1, p_2, \ldots, p_{n'}, \ldots, p_{|\mathcal{N}|}\right)}{\Delta} \tag{B.1}$$

for some very small Δ. This requires evaluating the city-wide housing demand function $2 \cdot |\mathcal{N}|$ times. Each of these evaluations requires solving an inner fixed point to make average incomes implied by location choices consistent with the average incomes that *motivate* those choices. For large cities with thousands of neighbourhoods, solving this fixed point can take over one second, and therefore evaluating the full Jacobian can take one to two hours. This is acceptable if the Jacobian only needs to be computed once, which is the case in the model inversion, but when solving for counterfactual scenarios computing this elasticity needs to be done many times in order to evaluate convergence criteria. As a result, the feasibility of this whole exercise depends on computing this Jacobian more quickly.

One may also consider using automatic differentiation software, such as ForwardDiff.jl or Zygote.jl. This is unfortunately not possible in my setting, as evaluation of the housing demand function requires solving the inner fixed point with endogenous amenities to make sure that they are consistent with households' choices. Since this is the product of a 'while' loop, automatic differentiation packages are not able to invert the exact sequence of arithmetic operations involved in this procedure from the compiled code, and therefore unable to differentiate them.