THE UNIVERSITY OF CHICAGO

ESSAYS ON SPATIAL AND URBAN ECONOMICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY
SHIJIAN YANG

CHICAGO, ILLINOIS
JUNE 2026

Copyright 2026 by Shijian Yang

All Rights Reserved

ABSTRACT

This dissertation examines the measurement, economic effects, and optimal design of land use regulations within a city.

In the first part of the dissertation, I examine how to quantify land use regulations. Land use regulations are difficult to measure because they consist of bundles of statutory rules whose effects on housing supply depend on local market conditions. To address this challenge, I construct a wedge-based measure of regulation that is recovered using data on construction costs and property values for Chicago and the surrounding Cook County.

In the second part, I model the economic effects of neighborhood-level regulations using a general equilibrium quantitative spatial model of a city with productivity and amenity externalities. This framework allows me to account for how urban structure and population distribution change with land use regulations. I calibrate the model to Cook County using commuting, income, and transit data.

In the third part, I study the welfare-maximizing configuration of land use regulations within this general-equilibrium framework and compare it to existing regulations. I numerically solve for the optimal allocation of regulations across neighborhoods in Cook County, using automatic differentiation to make this high-dimensional policy problem computationally tractable. I find that the optimal policy delivers substantial welfare gains relative to both current regulations in Chicago and wholesale deregulation. Optimal regulations depend less on local amenities and productivity in isolation than on spatial linkages across locations, measured by residents' access to jobs and firms' access to workers.

Taken together, this dissertation provides an integrated framework for land use policy that quantifies regulations, models their economic effects, and characterizes welfare-maximizing policies. The automatic-differentiation methods explored herein also offer a promising approach to a broader class of high-dimensional optimal policy problems.

For my family and friends

ACKNOWLEDGMENTS

This dissertation would not have been possible without the generosity of many people.

I am deeply grateful to my advisors, Esteban Rossi-Hansberg, Daniel McMillen, and Jonathan Dingel, for their guidance, patience, advice, and support.

I am also very grateful for the insightful feedback of Milena Almagro, Greg Kaplan, Rodrigo Adao, Erik Hurst, Chad Syverson, Kwok-Hao Lee, Xiaoyang Li, Chase Riley Abrams, Jordan Rosenthal-Kay, Zhiyu Fu, Xianglong Kong, Feng Lin, Nimisha Gupta, Andrew Koh, Yidi Wu, and numerous conference and seminar participants.

My research benefited immensely from institutional support. I thank the Becker Friedman Institute for Economics at the University of Chicago and the Fama-Miller Center for Research in Finance for access to CoreLogic data. I also thank the Singapore Social Science Research Council for support through its Graduate Research Fellowship.

I owe a profound debt of gratitude to my parents, Gao Yulu and Yang Mingchu, for their love, support, and encouragement throughout my life. Their sacrifices and care have made my achievements possible.

To my fiancée, Xuening: thank you for your love and unwavering belief in me throughout this journey. You bring me profound joy. After spending so long geographically apart, I am thrilled that we are beginning our life together. You remind me that there is life outside economics, and you are a constant source of strength and motivation every day. I could not have done this without you.

My family, friends, colleagues, and mentors have made me better in countless ways, and their support helped me do the best research I could. All errors, of course, remain my own.

TABLE OF CONTENTS

ACKNOWLEDGMENTS v

LIST OF FIGURES viii

LIST OF TABLES ix

1 INTRODUCTION 1

1.1 Contributions and overview 3

1.2 Related literature 6

2 MEASURING LAND USE REGULATION STRINGENCY 8

2.1 Land use regulations across the United States and common measures of regulation 8

2.2 A model of housing developers and land use regulatory wedges 10

2.3 Parcel-level supply and land values 13

2.4 Aggregation to neighborhood supply 15

2.5 Interpreting wedge τ as land use regulations 17

2.6 Recovering plot-level τ_ω and ZIP code land use stringency T_{Ri} 20

2.6.1 Improvement values 23

2.6.2 Property values 25

2.6.3 ZIP code-level aggregate wedges 28

2.6.4 Testing wedges against zoning in Chicago 29

3 A MODEL OF REGULATIONS ACROSS THE CITY 32

3.1 Quantitative Urban Model 33

3.1.1 Individual Utility 33

3.1.2 Residence-workplace choice probabilities 35

3.1.3 Production 35

3.1.4 Real Estate 36

3.1.5 Adding up conditions 38

3.1.6 Equilibrium 38

3.2 Model Calibration . . . 41
3.3 Residents, workers, wages, commuting . . . 42
3.4 Estimating commuting elasticity . . . 44
3.5 Local neighborhood characteristics . . . 45

4 COMPUTING OPTIMAL LAND USE POLICIES . . . 52
4.1 Solving for equilibrium for a given level of land use regulations . . . 53
4.2 Optimal land use regulation problem . . . 53
4.2.1 Land use regulations are not first-best policy for correcting urban externalities, and admit no analytical solution . . . 54
4.3 Finding welfare-maximizing land use regulations using automatic differentiation 56
4.3.1 Implementation . . . 60
4.4 Optimal land use policies in Cook County . . . 62
4.4.1 No-externality case . . . 62
4.4.2 Optimal regulations under urban externalities . . . 64
4.4.3 The effect of regulations on local housing markets . . . 73
4.4.4 What determines optimal land use regulations? . . . 76

5 CONCLUSION . . . 79

A APPENDIX . . . 90
A.1 Proof that land use regulations are not first-best policies . . . 90
A.2 Wedge τ is observationally equivalent to a quantity limit . . . 95
A.3 Wedge τ corresponds to a shadow price of extra housing in the presence of a quantity limit . . . 97
A.4 Deriving resident and firm market access from commuting probabilities . . . 98
A.5 Automatic differentiation . . . 101
A.6 Implementation details for AD-based optimization . . . 103
A.6.1 Quasi-Newtonian approximation of the Hessian . . . 103
A.6.2 Optimization under constraints on T_{Ri} . . . 105
A.7 Optimal policies under $\phi_{amen} = 0.15, \phi_{prod} = 0.07$ . . . 107

LIST OF FIGURES

2.1 Estimated improvement share of value for single-family and multi-family housing, using RSMeans construction costs and tax assessment property values. The red line denotes the chosen unrestricted improvement share of $\gamma = 0.85$ 26

2.2 ZIP-code aggregate regulatory wedge T_{Ri} across Cook County ZIP code tabulation areas . Single family improvement shares from Davis et al. (2021), multi-family shares from RSMeans construction cost data and tax assessment values. 29

3.1 Composite productivities and amenities across ZCTAs in Cook County. Productivities are on the 10^6 scale as I denominate real estate in millions of square feet. Gaps in the map are forest preserves not included in ZCTA shapefiles. 48

3.2 Resident market access (left) and firm market access (right) across ZIP codes in Cook County. 50

4.1 Optimal T_R^* across different amenity and productivity spillovers ϕ_{amen}, ϕ_{prod}. . 65

4.2 Optimal land use regulations T_{Ri} across space, $\phi_{amen} = 0.15, \phi_{prod} = 0.07$ . . . 69

4.3 Worker and resident population changes, optimal regulations T_R^* vs. current regulations T_R^0 70

4.4 Worker and resident population changes, T_R^* vs. uniform T_R keeping aggregate city-level regulations constant 72

4.5 Changes in neighborhood regulatory wedge and housing quantity (%) under optimal regulations T_R^*, relative to baseline T_R^0. 74

4.6 Spatial changes (%) in housing prices and housing values under optimal regulations T_R^*, relative to baseline T_R^0. 74

4.7 Aggregate outcomes under complete deregulation and optimal regulation, % change relative to baseline T_R^0. 75

A.1 Illustration of a computational graph for $f(x, y) = \sin(xy) + x^2$ 101

A.2 Optimal land use regulations T_{Ri} across space, $\phi_{amen} = 0.15, \phi_{prod} = 0.07$ . . . 107

LIST OF TABLES

2.1 Summary statistics for property size, estimated property values, estimated construction costs from RSMeans, and implied improvement shares for Cook County. Standard deviations are in brackets. 25

2.2 Regression of τ_ω on FAR and zoning codes fixed effects 30

3.1 Parameter Values and Sources 41

3.2 Commuting elasticity estimates 44

4.1 Welfare, average regulation T_{Ri} across locations, average wages, and per-person land-rents across different scenarios 67

4.2 Regression of log of ZIP-code optimal regulation T^*_{Ri} on amenities, productivities, resident MA, and firm MA 77

CHAPTER 1

INTRODUCTION

Land use regulations are laws and rules that govern the development of real estate to influence housing supply and restrict housing density and quantity. They are implemented by local governments and take many forms, including height and floor-area ratio limits, minimum lot sizes, discretionary approvals, and development fees. A large literature (Ganong and Shoag, 2017; Glaeser and Gyourko, 2018; Hsieh and Moretti, 2019; Herkenhoff et al., 2018) documents stringent land use regulations that suppress housing supply and raise housing prices in the most productive American cities, such as New York and San Francisco, and estimates income gains of 5–10% from easing these constraints. These facts highlight the need to examine the reasons for land use regulation.

Urban externalities are often cited as a rationale for land use regulations. Restrictive land use policy can address negative amenity externalities arising from congestion and overcrowding in urban neighborhoods (Brueckner, 2007; Duranton and Puga, 2015). Land use policy can also encourage localized productivity and knowledge spillovers (Ciccone and Hall, 1996; Rossi-Hansberg, 2004; Combes et al., 2015) by concentrating people in or near productive locations such as downtowns. By shaping the distribution of people across space, land use regulations can internalize productivity and amenity externalities arising from the concentration of people within the city.

Existing evaluations of land use regulation typically benchmark current rules against *predefined* policy counterfactuals rather than *optimal*, welfare-maximizing policy. Common exercises relax regulations to a historical baseline or remove them wholesale (Hsieh and Moretti, 2019; Herkenhoff et al., 2018; Duranton and Puga, 2023; Babalievsky et al., 2023; Mei, 2022; Macek, 2024). For example, Babalievsky et al. (2023) calculates a 2.9% increase in welfare from deregulation. Given the stringency of regulation in many jurisdictions, almost

all chosen deregulations yield sizable gains. Yet these comparisons are not designed to pin down the welfare-maximizing spatial pattern of regulation across locations. Recent work on optimal spatial policy (Rossi-Hansberg et al., 2025; Fajgelbaum and Gaubert, 2020) emphasizes correcting amenity and productivity externalities with location-based transfers.

In this dissertation, I develop an integrated framework for land use regulation that spans (i) measuring the stringency of existing land use regulations, (ii) modeling the economic effects of regulations in general equilibrium, and (iii) computing welfare-maximizing regulations. First, I construct a price wedge measure of land use regulation stringency that can be recovered using real estate data on property values and construction costs. After that, I embed the regulatory price wedges into a general equilibrium urban model to account for the economic effects of land use regulations. Productivity and amenity externalities from the concentration of residents and workers create a role for regulation. I show theoretically that land use regulations cannot perfectly correct urban externalities to achieve first-best outcomes, in contrast to the location- and workplace-specific income transfers proposed in the optimal space-based policy literature. I propose a numerical procedure that computes neighborhood-specific optimal regulations by solving a nested fixed-point problem with gradients obtained via automatic differentiation. Calibrating the model for the city of Chicago and Cook County, the optimal regulation raises welfare by 2% relative to current regulations and differs markedly from commonly proposed no-regulation benchmarks.

I choose Chicago as my empirical setting because its urban structure and physical geography make it well suited to quantifying land use regulations and to studying optimal policy. Chicago is a commonly used empirical example in urban economics (McDonald, 2009; Ahlfeldt and McMillen, 2018; Shertzer et al., 2018). Its relatively flat, topographically uniform landscape makes the assumption of a common housing production function across space more plausible. Chicago is also a useful setting for studying optimal land use regulation rather than simply the gains from broad deregulation. Unlike places such as New York or the

San Francisco Bay Area, where regulations are often viewed as highly restrictive, Chicago is better characterized as a moderately regulated city. This makes it a more informative setting for asking what welfare-maximizing regulation should look like, rather than whether regulation should simply be relaxed across the board. Consistent with this motivation, the counterfactuals in Chapter 4 show that naive wholesale deregulation generates only modest welfare gains relative to the calculated welfare-maximizing policy.

1.1 Contributions and overview

I first develop an approach to quantifying and modeling regulations, which is the focus of Chapter 2. For my purpose of considering optimal policy within a city, I use a price wedge measure (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Babalievsky et al., 2023) of regulatory stringency to quantify land use regulations. Local governments impose a variety of land use regulations at every location; floor-area ratio limits, setback requirements, and development fees exist simultaneously on a single property. The way these rules interact and the extent to which any single regulation binds is complicated and hard to isolate. The fundamental aim of land use regulation is to reduce the amount of housing built per unit area of land. I measure land use regulation using a price wedge on housing. The price wedge is a virtual tax on the housing developer's output price that rationalizes the observed quantity of housing given prices. Intuitively, a developer who builds very little in a high-price area is inferred to face a large regulatory wedge due to stringent land use regulations. Under some standard assumptions on the developer's production function, the price wedge-based measure of regulation for a property is proportional to the ratio of construction cost to total property value. The price wedge measure of regulation lends itself to aggregation to neighborhood-level measures of land use stringency that can be incorporated into a general equilibrium model.

The empirical setting of my paper is the city of Chicago and the surrounding suburbs in Cook County, Illinois. To construct the price wedge measure of regulation, I estimate construction costs using RSMeans building cost data and property values using CoreLogic data on property characteristics and sales. I find that the wedge measure of regulation accurately reflects features of the Chicago zoning code and matches the general perception of land use regulatory stringency across Cook County.

To model the economic effects of land regulations in Chapter 3, I incorporate my housing-plot level problem of housing developers facing regulations into a general equilibrium quantitative urban model of a city. The model features workers who make residential and commuting choices across neighborhoods, supplying labor to profit-maximizing firms. Price-taking housing developers face regulations in the form of price wedges and supply housing in each neighborhood, given local price wedges. The model features urban externalities in the form of productivity and amenity spillovers—the concentration of workers and residents in a location can have either positive or negative effects on the productivity and amenity of all individuals in that location.

I calibrate the quantitative urban model to Cook County using data on commuting, wages, and real estate from various sources. By inverting the model, I recover amenities, productivities, and commuting elasticities for Cook County.

The presence of urban externalities in my model of the city introduces the scope for governmental intervention, whether in the form of land use regulations or other policies. In Chapter 4, I consider the problem of a policymaker who maximizes the average expected utility of all individuals in the city by setting land use regulations across neighborhoods in the city. I show that land use regulation cannot achieve first best outcomes that perfectly correct for urban externalities, in contrast to residence- and workplace-based income transfers proposed in the optimal space-based policy literature (Fajgelbaum and Gaubert, 2020; Rossi-Hansberg et al., 2025). This is because urban externalities are generated by people, while land

use regulations target housing. By the "principle of targeting" (Bhagwati and Ramaswami, 1963; Johnson, 1965), first-best policies should target the source of the externality (people) rather than a related good (housing). This implies that land use regulations are not first-best policies.

While land use regulations cannot perfectly correct urban externalities, they are the most common policy used by local governments to influence the distribution of people across space and address urban externalities. Common policy rationales for restrictive land use policies include the prevention of overcrowding and the preservation of local residential amenities. Given the absence of analytical expressions for optimal regulations, I solve numerically for the welfare-maximizing land use regulations across neighborhoods in Cook County.

Given the high dimensionality of the optimal land use policy problem, I use automatic differentiation methods to obtain welfare gradients to solve numerically for welfare-maximizing policies. Optimal policy problems are more commonly solved with Rust (1987) nested fixed-point (NXFP) methods with first-differencing. NXFP methods are computationally infeasible in my setting because computing gradients of the objective function and general equilibrium system with respect to high-dimensional land use policy requires solving for the general equilibrium thousands of times for *each* update of the policy vector. Instead, I obtain gradients via automatic differentiation, a procedure most commonly-used in machine learning and optimal control. By decomposing the general equilibrium system of equations into a series of elementary operations such as addition and exponentiation, a computational graph of the equilibrium system is formed. Automatic differentiation applies the chain rule along the computation graph of the equilibrium to obtain gradients used to update policy guesses in a computationally efficient manner, allowing me to numerically solve for optimal policies in a feasible amount of time.

I compute welfare-maximizing land use policies across neighborhoods in Cook County for a range of productivity and amenity spillover strengths. Relative to current regulations

in Cook County, eliminating all regulations raises welfare by about 0.1%-0.5%, whereas the computed optimal policy delivers more significant welfare gains of about 2%. The spatial pattern of the optimal land use policy regulations depends on the direction of amenity spillovers: when amenities rise with population density, land use regulation should be more permissive in high-amenity neighborhoods with strong access to high-wage jobs; if amenities fall with density, the prescription reverses and more stringent regulations in those neighborhoods are warranted to curb congestion and crowding. Despite the differences in *relative* spatial patterns of optimal land use regulation across neighborhoods and spillover scenarios, across all scenarios optimal policy entails a reduction in the aggregate *level* of land use regulations and the movement of residents from southern to northern Cook County due to the deregulation of highly-regulated suburbs.

1.2 Related literature

This dissertation adds to the large literature on the welfare effects of land use regulation and housing supply restrictions (Hsieh and Moretti, 2019; Herkenhoff et al., 2018; Duranton and Puga, 2023; Babalievsky et al., 2023). Existing work compares the current regulations to fixed counterfactuals (e.g., full deregulation). In contrast, I compute the welfare-maximizing pattern of land use regulation across locations and compare it to current policy.

In Chapter 2, I follow a model-based approach that uses the land and improvement share of property value as a regulatory wedge (Brueckner et al., 2017; Tan et al., 2020; Brueckner and Singh, 2020; Babalievsky et al., 2023). I implement a simplified version of Babalievsky et al. (2023)'s approach to constructing regulatory wedges from improvement shares and justify the Cobb–Douglas structure and the benchmark unrestricted improvement share using evidence from the housing production function literature (Epple et al., 2010; Ahlfeldt and McMillen, 2014).

This wedge-based measure captures all housing supply distortions induced by land use regulations using observed market outcomes. Relative to survey-based indices such as the Wharton Residential Land Use Regulation Index (Gyourko et al., 2008, 2021), the outcome-based wedge avoids survey attrition and coverage concerns. Finally, the wedge measure microfounds neighborhood housing supply under regulations in a theoretically consistent way and leads naturally the quantitative spatial model of the city developed in Chapter 3.

In Chapter 4, I add to the literature on optimal spatial policies within quantitative spatial models (QSM). Analytical characterizations of first-best spatial policy clarify how location-based income transfers can internalize productivity and amenity spillovers (Rossi-Hansberg, 2004; Fajgelbaum and Gaubert, 2020; Rossi-Hansberg et al., 2025). This dissertation considers the optimality of commonly observed land use regulations, moving away from analytically tractable income transfers. I use Fajgelbaum and Gaubert (2020)'s optimal spatial policy framework to show that land use regulations cannot perfectly internalize urban externalities. Allen et al. (2016) considers optimal land use in the form of single-use zoning within a QSM and conducts deregulation exercises with local perturbations to current regulations; I go further and solve numerically for globally optimal regulations.

Methodologically, I use automatic differentiation (AD) to obtain exact policy gradients and make gradient search in my high-dimensional optimal regulation problem computationally tractable. While standard in machine learning, AD is uncommon in economics; applications include computing exact derivatives of DSGE likelihoods (Bastani and Guerrieri, 2008) and solving nonlinear heterogeneous-agent models (Boehl, 2023). My work is part of the very recent literature on computing optimal policies using AD and machine learning methods, such as Wang et al. (2024) who compute optimal tariffs in a quantitative trade model using AD within a Rust (1987) nested fixed-point framework, and Sargent and Yang (2025a) who compute optimal money supply paths in a macro model using AD.

CHAPTER 2

MEASURING LAND USE REGULATION STRINGENCY

In this chapter, I quantify the stringency of land use regulations on individual properties using a price wedge-based measure. The wedge can be recovered from property characteristics and sales data. The price wedge measure of regulation lends itself to aggregation at the neighborhood and city levels. This yields a microfounded notion of neighborhood-level regulatory wedges that affect neighborhood housing supply and allows me to compute optimal neighborhood-level housing regulations in my general-equilibrium model of a city in Chapter 3.

2.1 Land use regulations across the United States and common measures of regulation

Measuring land use regulations is difficult because land use policy takes many shapes and forms. In the United States, land use regulation is set primarily at the local level by local municipal governments. Local governments adopt zoning ordinances and development codes, and establish planning and permitting bodies to regulate land use. Differing political and fiscal considerations across locations mean that the objects of land use policy and the policies used to control land use vary substantially across locations.

Common tools for regulating land use include single-use zoning (single-family, multifamily, commercial, industrial), density and quantity controls (floor-area ratio limits, height limits, minimum lot sizes), form and design standards (parking and yard setbacks, parking minimums). Governments also control land use via procedural and fiscal levers such as site-plan and environmental reviews, discretionary approvals and variances, and permit fees. Land use policy in practice is a bundle of regulations and processes that together determine

the cost and feasibility of real-estate development. (Glaeser and Gyourko, 2018; Gyourko and Molloy, 2015)

The most commonly-used measure of land use regulations is the survey-based Wharton Residential Land Use Regulatory Index (Gyourko et al., 2008, 2021) . The Wharton index is compiled from responses to a city-manager survey about local land use regulatory procedures, and is only available at the core-based statistical area (CBSA) level. Hsieh and Moretti (2019) and Baum-Snow and Han (2024) perform deregulation exercises by regressing housing-supply elasticities on the Wharton index and then plugging in alternative index values. This approach of using the Wharton index to adjust supply elasticities can be used for deregulation counterfactuals at the city/CBSA level. Given that I study regulations across neighborhoods within a city and want to consider optimal policy, I cannot use the Wharton index in my setting.

Researchers sometimes quantify regulation using a single provision of the zoning code. The floor-area ratio (FAR) limit, which can be thought of conceptually as a quantity limit on the amount of building given land size, has been frequently used as a measure of land use regulations, in works such as Brueckner and Singh (2020), Tan et al. (2020), and Moon (2019). Focusing on a single type of regulation is appropriate when the research question is about changing that specific rule (e.g., relaxing minimum-lot-size requirements in Macek (2024)). In Cook County, regulatory stringency cannot be easily represented by a single aspect of the local zoning code. Any given lot must simultaneously satisfy FAR caps, height limits, minimum lot sizes, and front- and rear-setback requirements, among others. Treating just one of those constraints as "the" measure of regulation will misstate the true regulatory stringency.

Commonly-used survey-based measures of regulation are geographically coarse, and single policy measures of regulation cannot reflect the totality of complex land use regulations. For my purpose of considering optimal land use policy in a general equilibrium model, I use a

price wedge measure of regulation simplified from the setting of Babalievsky et al. (2023). The price wedge measure of regulation is derived from a model of real estate development where developers face a price wedge on their output due to land use regulations. The wedge can be recovered from property characteristics and sales data.

2.2 A model of housing developers and land use regulatory wedges

To measure land use regulations, I start from a land parcel/property-level model of housing production by price-taking, atomistic real-estate developers, who produce real estate using land and improvements. Improvements, denominated in dollars (the numeraire), are a composite input that includes all non-land inputs used to produce real estate, such as labor, materials, and machinery. Splitting housing production into land and improvements is a standard assumption in the housing production function literature such as Epple et al. (2010) and Combes et al. (2021).

At each land plot ω, the developer faces land use regulation that reduces the amount of housing built in equilibrium. These regulations take the form of a price wedge τ_ω. The price wedge at each plot is a virtual tax on housing prices facing the developer that rationalizes observed quantities of building under complex land use regulations.

Each plot ω has fixed land size l_ω and is owned by atomistic landlords who cannot influence prices. There is an infinite mass of atomistic, price-taking real-estate developers who can potentially build at every plot. Land is sold to price-taking developers who choose improvements x_ω, denominated in the numeraire, and pay the per-unit market price for land q_ω. Land rents $q_\omega l_\omega$ adjust so that developers earn zero profits in equilibrium. The real estate production function at plot ω is $h(l_\omega, x_\omega; z_\omega)$, where z_ω is plot-level construction productivity. Land is a fixed factor of production and developers only choose their level of improvements x_ω, given per-unit price of real estate p_ω. I consider first a setting where there

are no regulations or distortions limiting the developer's choice of quantities. The developer's profit maximization problem, in an unregulated world, is thus

$$\max_{x_\omega} \; p_\omega \; h(x_\omega, l_\omega; z_\omega) - x_\omega - q_\omega l_\omega \, .$$

Since l_ω is fixed from the developer's perspective, the profit-maximizing first-order condition for the developer's choice of improvements x_ω is given by:

$$1 \; = \; p_\omega \, \frac{\partial h(x_\omega, l_\omega; z_\omega)}{\partial x_\omega}.$$

It is commonly assumed that the housing production function is a constant returns to scale Cobb-Douglas function; $h(x, l; z) = z \, x^\gamma l^{\,1-\gamma}$ with $0 < \gamma < 1$. This is not just a convenient assumption, but rather an evidence-based consensus in the housing production literature. Ahlfeldt and McMillen (2014), using a large land and property sales dataset spanning the United States, and Epple et al. (2010) , using data from Allegheny county, Pennsylvania, both show that the elasticity of substitution between land and improvements is very close to 1. Both studies also find that housing production function has close-to constant elasticity of substitution between land and improvements; second order exponents in a general translog production function specification are very close to zero. A constant elasticity of substitution production function, with the elasticity of substitution being close to 1, justifies the usage of a Cobb-Douglas production function for real estate.[1]

Both Epple et al. (2010) and Ahlfeldt and McMillen (2014) estimate the Cobb-Douglas exponent on improvements in the production function, γ, to be approximately 0.85. If this is the case, then in an unrestricted world where the profit-maximizing developer does not

1. Combes et al. (2021), which uses French rather than American data, calculates a elasticity of substitution between land and improvements of less than 1, but still finds that Cobb-Douglas approximates the true housing production function reasonably well.

face any sort of land use regulation improvement values should be 85% of the sales price of real estate, with land value being the remaining 15% of house prices.

Clearly, this is not the case in reality. In locations generally thought to have high levels of land use regulation, improvement share (which is also the construction cost share, excluding land cost) is much lower than 85%, while land captures a much larger fraction of total value.

In the spirit of "wedge accounting" (Chari et al., 2007; Desmet and Rossi-Hansberg, 2013; Babalievsky et al., 2023), I model regulation as a reduced-form price wedge/virtual distortion τ that enters the developer's first-order condition so that the observed quantities of housing built, given prices, can be interpreted as the solution to the developer's problem with an additional output tax. The wedge τ is the housing analogue of an output wedge in business-cycle accounting: it depresses the marginal revenue product of improvements, tilts revenue shares from capital/improvements toward land, and scales down equilibrium quantities while remaining agnostic about the exact form of land use regulation (minimum lot sizes, FAR limits, setbacks, development fees etc.) that depresses housing quantities .

In a world with land use regulation, plot-level land use regulation enters the developer's problem as a wedge $\tau_\omega \in [0, 1]$ on the effective sales price, so the developer's problem for a parcel of size l_ω is

$$\max_{x_\omega} \ \tau_\omega \, p_\omega \, h(x_\omega, l_\omega; z_\omega) - x_\omega \, .$$

The first–order condition is

$$1 = \tau_\omega \, p_\omega \, \frac{\partial h(x_\omega, l_\omega; z_\omega)}{\partial x_\omega} \, . \tag{2.1}$$

The price wedge τ_ω reduces the marginal revenue product of improvements. A low value of τ_ω indicates binding and stringent land use regulations, while a higher value of τ_ω indicates looser regulatory stringency. $\tau_\omega = 1$ would indicate that there are distortions in the developer's problem, and thus no or minimal regulation. Rearranging equation (2.1) yields:

$$\tau_\omega \;=\; \frac{1}{p_\omega\,\partial h_\omega/\partial x_\omega} \;=\; \frac{1}{\frac{\partial h_\omega}{\partial x_\omega}\frac{x_\omega}{h_\omega}}\;\frac{x_\omega}{p_\omega\,h_\omega}\,. \tag{2.2}$$

$\frac{\partial h}{\partial x}\frac{x}{h}$ is the elasticity of housing output with respect to improvements, denoting how much more real estate is produced with an increase in improvements, holding land constant. Using the evidence from the housing production literature that the production function $h(x,l;z)$ takes a Cobb-Douglas form, I write $h(x,l;z) \equiv z\,x^{\gamma}l^{\,1-\gamma}$. Under this functional form of the production function, the output elasticity with respect to improvements is constant; $\frac{\partial h}{\partial x}\frac{x}{h} = \gamma$. We can then write τ_ω, the price wedge measure of regulation for plot ω, as

$$\tau_\omega \;= \frac{1}{\gamma}\;\frac{x_\omega}{p_\omega\,h_\omega} \;=\; \frac{1}{\text{output elasticity of improvements}} \times \frac{\text{improvement value}}{\text{total property value}}. \tag{2.3}$$

The plot-level wedge τ_ω that rationalizes low building quantities can be recovered using observed total property value and improvement cost data. Observed sales prices identify total property values $p_\omega\,h_\omega$, while the improvement values x_ω, which are defined as the cost of all inputs other than land, correspond to construction costs.

2.3 Parcel-level supply and land values

Solving the profit-maximization problem of a housing developer who faces price wedge τ_ω yields the following improvement choice x_ω and real-estate output h_ω, in terms of real-estate price p_ω and fixed parcel-specific land size l_ω:

$$x_\omega = \left(\gamma\,p_\omega\,\tau_\omega\,z_\omega\right)^{\frac{1}{1-\gamma}} l_\omega, \qquad h_\omega = \left(\gamma\,\tau_\omega\right)^{\frac{\gamma}{1-\gamma}} z_\omega^{\frac{1}{1-\gamma}}\,p_\omega^{\frac{\gamma}{1-\gamma}}\,l_\omega\,. \tag{2.4}$$

Improvement share of total value is $\frac{x_\omega}{p_\omega h_\omega} = \gamma\tau_\omega$. Competition among developers drives profits to zero and land is the residual claimant to the total property value ph. Land value

ql is

$$q_\omega l_\omega = p_\omega h_\omega - x_\omega = (1 - \gamma\tau_\omega) p_\omega h_\omega \,.$$

Wedge τ_ω not only reduces the marginal revenue product of improvements and thus real estate output, but also increases the landlord's share of total housing value.

A wedge that raises land's share of real-estate value may seem counterintuitive at first. In fact, it captures how land use regulation is capitalized into land prices. Consider a high-demand suburb with binding floorspace limits; i.e. low-density housing in an affluent, highly regulated area. Decomposing property value ph into structure and land components, $p_\omega h_\omega = x_\omega + q_\omega l_\omega$, the binding floorspace constraint keeps x_ω (the value of the small structure) from rising with demand, while scarcity rents are capitalized into land value $q_\omega l_\omega$. Property values at regulation-constrained locations are high not because small buildings in those locations are costly, but because the right to build is scarce. Land use regulations push up land price $q_\omega l_\omega$ and the land share of total property value $\frac{q_\omega l_\omega}{p_\omega h_\omega}$.

The wedge τ_ω is a distortion that rationalizes the amount of housing built per unit of land, given observed prices. Modeling regulations as wedge τ_ω does not mean that regulations must take the form of a tax on housing output. In practice, land use policies (e.g. height/FAR limits and minimum lot sizes) reduce building quantities and create scarcity rents capitalized into land values. The wedge τ_ω replicates the observed effects on building quantity, improvement choices, and land values of real-world land use policies. I illustrate the *observational equivalence* between a price wedge τ and a binding quantity limit $\bar{h}$ (a density/FAR cap) in Appendix A.2. For any binding $\bar{h}$, one can choose τ so that the profit-maximizing developer's optimal choice equals $\bar{h}$. The wedge τ and the quantity limit $\bar{h}$ generate the same housing quantities and the same capitalization into land values.

2.4 Aggregation to neighborhood supply

Let i index a neighborhood (or any other unit of geography) composed of parcels $\omega \in i$. Neighborhood-level housing supply and total improvement values are the sum of plot-level supply and plot-level improvements in equation (2.4):

$$\sum_{\omega \in i} h_\omega = \sum_{\omega \in i} \left(\gamma\, \tau_\omega\right)^{\frac{\gamma}{1-\gamma}} z_\omega^{\frac{1}{1-\gamma}}\, p_\omega^{\frac{\gamma}{1-\gamma}}\, l_\omega \tag{2.5}$$

$$\sum_{\omega \in i} x_\omega = \sum_{\omega \in i} \left(\gamma\, p_\omega\, \tau_\omega\, z_\omega\right)^{\frac{1}{1-\gamma}}\, l_\omega \tag{2.6}$$

Neighborhood-level housing supply, which is the sum of individual developers' plot-level supply, can be represented as the outcome chosen by a single *representative* neighborhood developer. This representative developer faces an *aggregate* residential land use wedge T_{Ri}, uses neighborhood-level improvements X_i on aggregate land L_i with aggregate efficiency Z_i, and produces total housing H_i. Suppose that:

$$T_{Ri} \;=\; \left(\frac{\sum_{\omega \in i} x_\omega / \tau_\omega}{\sum_{\omega \in i} x_\omega / \tau_\omega^{\frac{1}{1-\gamma}}}\right)^{\frac{1-\gamma}{\gamma}}, \tag{2.7}$$

$$L_i \;=\; \sum_{\omega \in i} l_\omega, \tag{2.8}$$

$$Z_i \;=\; \left(\frac{1}{L_i} \sum_{\omega \in i} z_\omega^{\frac{1}{1-\gamma}}\, l_\omega\right)^{1-\gamma}. \tag{2.9}$$

A representative developer with Cobb–Douglas production in neighborhood-level land and improvements, $Z_i X_i^\gamma L_i^{1-\gamma}$, and facing neighborhood-level wedge T_{Ri}, solves

$$\max_{X_i} \; T_{Ri}\, p_i\, Z_i\, X_i^\gamma L_i^{1-\gamma} \;-\; X_i.$$

The representative developer then chooses neighborhood improvement quantities X_i and housing supply H_i:

$$X_i = \left[\gamma\, T_{Ri} Z_i p_i\right]^{\frac{1}{1-\gamma}} L_i, \tag{2.10}$$

$$H_i = \left[\gamma\, T_{Ri}\right]^{\frac{\gamma}{1-\gamma}} Z_i^{\frac{1}{1-\gamma}} L_i\, p_i^{\frac{\gamma}{1-\gamma}}. \tag{2.11}$$

We can verify algebraically that the definitions of T_{Ri}, L_i, and Z_i in equations (2.7), (2.8), and (2.9) yield neighborhood-level housing quantities and improvement choices H_i and X_i that match the aggregation of individual plot-level choices in equations (2.5) and (2.6); specifically, $X_i = \sum_{\omega \in i} x_\omega$ and $H_i = \sum_{\omega \in i} h_\omega$. In this sense, T_{Ri} is a neighborhood-level wedge that reproduces the same aggregate improvement choice and housing supply as the decentralized sum of plot-by-plot choices. We can therefore write neighborhood-level housing supply H_i in terms of the neighborhood wedge T_{Ri} and neighborhood real-estate price p_i:

$$H_i = \left[\gamma\, T_{Ri}\right]^{\frac{\gamma}{1-\gamma}} Z_i^{\frac{1}{1-\gamma}} L_i\, p_i^{\frac{\gamma}{1-\gamma}}.$$

The aggregation result provides a microfoundation for the neighborhood-level housing supply function by deriving it directly from the plot-level problem of housing developers facing land-use regulations. It therefore gives neighborhood-level wedges a clear theoretical interpretation as measures of regulatory stringency. This contrasts with survey-based measures such as the Wharton Residential Land Use Regulatory Index, which are not derived from an explicit model of housing supply and are typically used in an atheoretical manner, for example by relating local housing supply elasticities to the index in reduced-form regressions.

This microfounded neighborhood-level housing supply function allows me to study housing supply under land-use regulations within a general equilibrium framework. I incorporate

aggregate neighborhood-level supply functions directly into the housing supply equations (3.8) of my general equilibrium quantitative spatial model of Cook County in Chapter 3.

The aggregation result thus also allows me to formulate and solve the optimal land use regulation problem at the neighborhood level rather than at the parcel level in Chapter 4. Solving directly for welfare-maximizing τ_ω with a general equilibrium model for each of the 1.1 million land parcels in Cook County is computationally infeasible because the policy vector would be extremely high-dimensional.

2.5 Interpreting wedge τ as land use regulations

The wedge τ is a reduced-form distortion that captures the gap between the price of housing and its marginal cost of production. A low τ means that the price is significantly above marginal cost because supply cannot expand freely. In the land use context, this reflects binding regulatory limits on density or quantity that act like an implicit tax on construction by preventing additional development even when it would be profitable.

τ can be seen as reflecting the shadow price of these constraints. Appendix A.3 shows the mathematical relationship between the shadow price of a binding quantity constraint and τ. A higher shadow price of a binding quantity constraint corresponds to a lower τ representing more restrictive land use regulations.

It is important to emphasize that τ is not a direct measure of the written regulations but rather how restrictive the land use regulations are, given equilibrium prices. Identical formal regulations, such as the same FAR ratio quantity limit, can lead to very different τ values across neighborhoods. A location with a given quantity limit may exhibit a large τ if housing prices are high due to high demand, whereas the same quantity limit in a low-demand area would result in a negligible τ if the constraint does not bind.

For intuition, consider a hypothetical rule that only single-family housing at a fixed low density can be built anywhere in Chicago. On paper, this policy is uniform. In practice, it would barely bind in some low-income parts of the South Side but would be extremely restrictive in downtown Chicago, where high-rise multifamily buildings are currently observed. In downtown, the implied wedge τ would be large, because dense redevelopment is profitable at current prices but blocked by regulation. The wedge τ therefore captures the *stringency* of regulation, not just its legal text. It can also reflect informal or negotiated restrictions, such as political bargaining or discretionary approvals that limit quantities. In general, τ reflects any force that keeps observed housing quantities low relative to what prices would otherwise justify.

The regulatory wedge τ is identified from variation in the improvement share of property value. This interpretation requires an identifying assumption: absent land use regulations, the housing production technology, and therefore the unconstrained improvement share of value, would not vary systematically across parcels and neighborhoods. Put differently, cross-neighborhood differences in observed improvement shares are interpreted as reflecting differences in the stringency of constraints on building, rather than large differences in the underlying technological relationship between land and structures. In principle, this assumption could fail if some neighborhoods faced systematically tighter natural supply constraints, such as steep slopes, irregular terrain, or other geographic barriers that limit buildability and raise the cost of expanding housing supply, as emphasized by Saiz (2010). However, such concerns are less salient in Cook County. Unlike cities (such as Los Angeles and San Francisco) where topography sharply limits development, Cook County is uniformly flat, so geographic housing supply constraints of this kind are unlikely to generate meaningful cross-neighborhood variation in improvement shares or housing supply in my setting. This makes Cook County a particularly suitable setting for interpreting low improvement shares as evidence of tighter land use regulation rather than as the consequence of heterogeneous

natural constraints on construction.

τ may capture distortions to building that do not arise directly from land use policy. These include market power in land and construction markets, as well as credit constraints that limit building (D'Amico and Alekseev, 2025). More generally, τ reflects the overall price-cost gap in housing markets, and that gap cannot be attributed exclusively to zoning (Glaeser and Gyourko, 2018). Empirically, however, local construction markets in much of the United States are highly fragmented and appear broadly competitive (Ahluwalia et al., 2022), so market power is unlikely to be a major driver of the gap between prices and marginal costs. I examine the universe of Chicago building permits issued in 2019 and find substantial developer fragmentation: the median ZIP code has 43 distinct housing developers, while the smallest number of developers in any ZIP code is 10. This pattern suggests limited market power in the local housing construction sector and indicates that there are no strong barriers preventing developer entry across neighborhoods within Cook County. It is therefore reasonable to attribute most cross-neighborhood distortions in housing markets to variation in land use stringency. The coexistence of locations with high and low τ within Cook County is also more consistent with differences in regulatory stringency than with persistent local market power. τ should therefore be interpreted as an upper bound on the effect of land use regulation: it captures the total elevation of price above marginal cost, of which land use regulations are typically a major component, but not the sole source.

τ reflects distortions at the intensive margin of building and is less informative about the extensive margin. The regulatory wedge is measured for properties that are already built, but it does not tell us why a given land parcel is developed or empty. If a policy prevents any development from happening on a parcel (e.g., green belts banning all construction at the city periphery (Koster, 2024)), then parcel supply is zero regardless of price. Similarly, τ cannot directly measure other extensive-margin distortions such as permitting delays, uncertain discretionary reviews, or the need to overcome community opposition to building.

While we can attribute a τ of zero to parcels where development does not happen at all, the wedge measure of regulation is not a useful concept when thinking about land use policy on the extensive margin.

The price wedge is informative when evaluated not just on new housing stock but also on *existing* housing stock. For cities like Chicago with large numbers of legacy buildings, a low τ signals that additional land use intensification (via renovation, redevelopment, infill, or additions) that would be privately profitable at prevailing prices is limited by constraints on development. In this case, τ_i reflects the unrealized value of land use intensification. The age of properties does not negate the information in τ, as it indicates whether existing regulations restrict further development relative to current willingness to pay for such improvements.

This interpretation is especially salient in high-price cities that nonetheless exhibit low densities across many neighborhoods (e.g., San Francisco). In such situations, low τ and low observed density indicate that, absent constraints, the market would redevelop lots to higher intensity. Recent reforms in California (Marantz et al., 2023), which relax land use policy by legalizing second housing units on existing single-family land parcels in the form of accessory dwelling units (ADUs), act directly to increase improvement intensity and raise τ_ω.

2.6 Recovering plot-level τ_ω and ZIP code land use stringency T_{Ri}

ZCTA-level residential land use regulatory wedge T_{Ri} is the aggregation of individual parcel-level wedges, using the aggregation results outlined in Section 2.4. While my model allows for land use regulation on commercial real estate, I do not conduct any policy experiments on commercial land use regulations due to a lack of data. I set commercial regulatory stringency $T_{Ci} = 1$ over all locations; differences in the stringency of commercial land use regulations instead get absorbed into the local commercial real estate construction productivity Z_{Ci}.

From equation (2.3), parcel-level regulatory wedge τ_ω can be recovered using the improvement share of property value:

$$\tau_\omega = \frac{1}{\gamma_R} \times \frac{\text{Improvement value}}{\text{Property value}}.$$

I assume a value of $\gamma_R = 0.85$ using estimates from Ahlfeldt and McMillen (2014) and Epple et al. (2010), meaning that without regulations improvement value should be 85% of total property value.

My data source for property characteristics and tax assessments for properties in Cook County comes from CoreLogic/Cotality [2]. CoreLogic/Cotality is a property data vendor that compiles and harmonizes parcel-level property characteristics, tax, and sales data from local government records, multiple listing service (MLS) feeds, insurance data, mortgage performance data, and other sources into near-nationwide files. From the CoreLogic Tax and Property Assessment data and the CoreLogic Owner Transfer data, I am able to obtain rich property characteristics, tax assessments, and sales history of residential and commercial properties in Cook County.

I obtain improvement share of total value from two sources. For single-family residences, I use Davis et al. (2021)'s published improvement share of value for single family residences. I construct improvement share of value for multi-family residences by obtaining structure value estimates from a construction industry cost provider, RSMeans, and inferring property values from tax assessment data.

My source for the improvement share of value for single-family residences comes from Davis et al. (2021). Using mandatory appraisal reports from 2012 to 2019 submitted to

2. The company previously named CoreLogic was renamed Cotality in March 2025. I retain the historical name "CoreLogic" throughout to refer to the pre-2025 CoreLogic-branded dataset names used in this paper. Cotality claims to have data for more than 99.9% of all properties in the United States. Generally, economists have treated its dataset as containing the census of properties in the USA.

Fannie Mae and Freddie Mac, the two large government-sponsored mortgage market intermediaries, Davis et al. (2021) recover land and improvement values and value shares for single-family homes at the ZIP code, census tract, and county levels across the United States. Property appraisals are conducted manually by professional appraisers with local market knowledge in the course of property and mortgage transactions. Appraisal reports also contain accurate estimates of the improvement and land values underlying the property; improvement values measure structure replacement costs, which are important for risk management. I take ZIP code tabulation area-level improvement shares for single-family housing directly from their data.

Of the 1,105,704 residential properties in my Cook County CoreLogic data, 17.2% are multifamily properties, accounting for 26.5% of total floorspace. I need to account for land use regulations on multifamily housing to get an accurate picture of regulations across Chicago, especially since the central business district and downtown areas of Chicago ("The Loop") contain almost no single-family housing while accounting for a significant portion of both residences and workplaces in Chicago.

I classify multifamily residences into two types: low- and medium-density residences, which are 1 to 3 stories tall, and high-density residences, which are more than 4 stories. To estimate the regulatory wedge for low- and medium-density multifamily residences, I use RSMeans construction cost data to estimate improvement values and tax-assessment data to measure total property value. A significant portion of low- and medium-density multifamily buildings have high improvement shares that imply minimal regulation ($\tau_\omega \approx 1$). RSMeans data does not provide reliable estimates for high-density multifamily properties of four stories or more. Given that taller and denser buildings generally face less stringent land use regulation, I assume that high-density residences do not face binding land use regulations and set $\tau_\omega = 1$ for buildings that are four or more stories tall.

2.6.1 Improvement values

Improvement value is the value of the non-land component of a property. For a newly built structure, this coincides with construction cost. For an existing structure, however, the relevant object is depreciated replacement cost: the current cost of constructing a comparable structure, net of physical depreciation. I therefore use the term improvement value to refer to depreciated replacement cost rather than the cost of a new property with the same characteristics.

To estimate improvement values, I use RSMeans Residential Cost Data, a building-cost dataset widely used in the construction, insurance, appraisal, and real estate industries. Davis et al. (2021) note that professional property appraisers commonly rely on RSMeans (or its competitor Marshall & Swift) as the underlying construction cost source and define structure value using depreciated replacement cost. RSMeans has also been used by economists such as Glaeser and Gyourko (2018), Duranton and Puga (2023), and Potter and Syverson (2025) to estimate construction costs.

RSMeans contains detailed unit-level and task-based costs for materials, labor, and equipment based on field research, structured surveys of manufacturers, suppliers, and contractors, and cost modeling. The RSMeans data contains both national average costs and localized city cost indices to adjust for differences in costs across locations. RSMeans supports construction cost estimation at both coarse and fine levels of detail, up to granular task pricing (e.g., comparing roofing systems such as 4-ply organic asphalt with insulated deck vs 3-ply coal-tar), enabling consistent translation from building specifications to physical cost estimates.

Since I do not observe task-level construction specifications for each property, I rely on RSMeans' standardized square-foot building models to impute costs. These provide per-square-foot cost estimates for a wide range of standardized residential property structures

at a high level of detail: 1-, 1.5-, bi-level, 2-, 3-, story structures across construction classes (economy, average, custom, luxury), exterior wall systems (e.g. wood, brick, stucco, masonry finishes), and basement finishes. The data also includes cost adjustments for common property features such as attached/detached garages and additional bathrooms.

I digitize the published RSMeans Residential Cost Data for a range of standardized building types. The CoreLogic dataset contains the building attributes needed to map each property to the closest RSMeans prototype, including square footage, number of stories, exterior finish, basement finish, garage type, and bathroom count. This mapping yields a parcel-level estimate of undepreciated replacement cost for the structure. I then convert replacement cost new into depreciated replacement cost, which I use as my estimate of improvement value x_ω for multifamily buildings. Table 2.1 reports summary statistics for the resulting estimates.

It is important to account for depreciation. RSMeans estimates the cost of constructing an equivalent new building today based on the physical characteristics observed in the CoreLogic data, such as square footage and building materials. Using this undepreciated estimate directly would overstate the value of the existing structure and therefore understate the regulatory price wedge. Property values are also shaped by age: holding other characteristics fixed, older buildings typically sell for less than newer ones because of physical deterioration, functional obsolescence, and lower remaining useful life. A depreciation adjustment to construction costs is therefore necessary to make the estimated improvement value comparable to the market value of the property.

The decomposition of property value into land and structure becomes less reliable for older buildings. For relatively old homes, total property value reflects not only the value of the land and the depreciated value of the existing structure, but also the option value of demolishing the current building and redeveloping the parcel. In a simulation exercise, Davis et al. (2021) show that this redevelopment option is small for relatively young houses: for

homes built within the previous 20 years, the option value accounts for less than 5 percent of total property value under their benchmark calibration. In that case, total property value is well approximated by the sum of land value and depreciated improvement value. Following this logic, I restrict my sample to buildings less than 20 years old when constructing regulatory wedges. I also apply a depreciation rate of 2.3 percent, following Harding et al. (2007) and Davis et al. (2021).

2.6.2 Property values

The estimation of τ_ω requires both improvement values and total property values. Sales values are, by definition, property values. I obtain my estimate of 2019 property values from assessment data from the Cook County Assessor, which is reported in CoreLogic. The Cook County Assessor's Office derives property values using an open-source LightGBM gradient-boosting machine-learning model from sales data. This follows a history of scandals surrounding unequal assessments, which led to significant reforms, including open-sourcing both code and data for the sales-data-driven automated valuation model in the tax assessment process, significantly reducing assessment regressivity and improving assessment accuracy (Berry, 2025).

	All properties	Single-family	Multi-family
Building sq footage	1878 (1343)	1641 (998)	3066 (2039)
Lot size (sqft)	6686 (6191)	7163 (6570)	4397 (2959)
Improvement value	238,568 (119396)	216,759 (87372)	347,847 (181955)
Property value	295,013 (329313)	282,263 (278276)	357,033 (504629)
Improvement share	0.76 (0.41)	0.74 (0.40)	.890 (0.45)
No. obs	1,023,662	853,352	170,310

Table 2.1: Summary statistics for property size, estimated property values, estimated construction costs from RSMeans, and implied improvement shares for Cook County. Standard deviations are in brackets.

Cook County is divided into 3 tax assessment zones, with assessments updated in a

three-year cycle. I use assessments from the 2019-2021 assessment cycle, which saw the implementation of the sales-data-driven valuation model. For a like-to-like comparison, I adjust the 2020 and 2021 assessed values to 2019 values using price indices constructed from a hedonic regression of sale prices of arm's-length transactions against property characteristics and sales year fixed effects. Compared to 2019, sales values were 8.6% higher in 2020 and 14.4% higher in 2021. I find that the assessed property values are a good predictor of property sales values, with a simple regression of arm's-length logged sales prices against logged assessed values yielding an R^2 of 0.6.

Table 2.1 reports assessed property values and estimated improvement values derived from the RSMeans cost data. I divide the estimated improvement value by property value to obtain an estimate of the improvement share of total value, which directly informs τ_ω given my choice of γ_R.

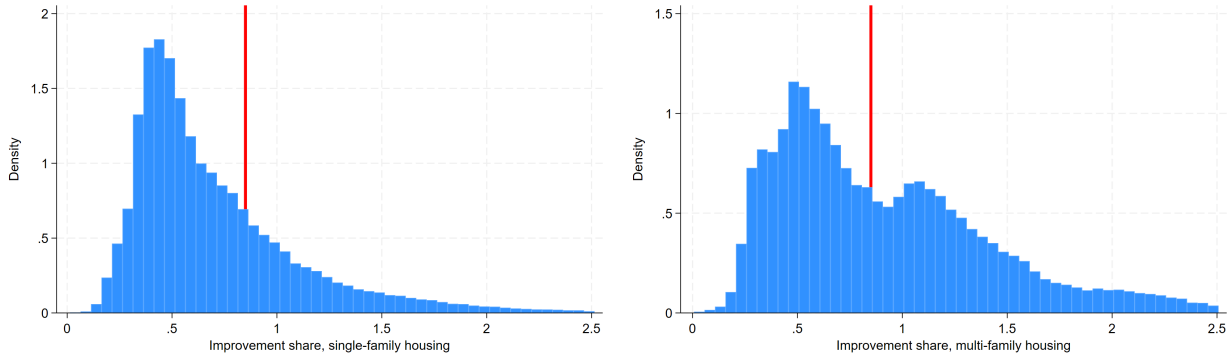

Figure 2.1: Estimated improvement share of value for single-family and multi-family housing, using RSMeans construction costs and tax assessment property values. The red line denotes the chosen unrestricted improvement share of $\gamma = 0.85$

Figure 2.1 plots the distribution of improvement share of value estimated using RSMeans construction costs and tax assessment property values for both single- and multi-family properties. 18.2% of single-family housing and 38.4% of multi-family housing have improvement shares larger than 1. An improvement share larger than 1 would imply both negative land values and a negative regulatory wedge. This would imply a virtual subsidy on construc-

tion. As discussed below, the large number of properties with calculated improvement shares larger than one is due to the limitations of the estimated construction costs.

Many properties in the South and Southwest sides of Chicago were built during the 1890–1929 construction boom and subsequently faced a long period of population loss and job decline after the 1960s due to deindustrialization and white flight (Pacyga, 2009). There is a significant aging stock of large structures in Chicago that are oversized for current market demand and prices. In fact, there are 3 ZIP codes (60636 West Englewood, 60472 Robbins, 60426 Harvey) where mean property values are below \$65880, the depreciated structural costs of a small, 30-year-old, 1000 square foot, basic lumber-sided house[3].

In these low-price neighborhoods, this legacy stock yields a mechanical mismatch: even after depreciation adjustments to construction costs estimated using RSMeans, the estimated structure costs can exceed the property's total value, implying an improvement share above one.[4].

In such properties, real estate regulations are not the binding margin; the persistence of large, older structures alongside very low prices is driven by weak demand and obsolescence. I treat parcels with improvement shares greater than my assumed value of $\gamma = 0.85$ as minimally regulated, with $\tau_\omega = 1$. There is a long left tail of both single- and multifamily housing with estimated improvement shares larger than one in my sample. It is for this reason that I use single-family improvement shares from Davis et al. (2021), which uses improvement values from professional property appraisals. From Davis et al. (2021)'s data, most South Side ZIP codes have improvement shares between 0.8 and 0.85, implying minimal regulation with my choice of $\gamma = 0.85$. There are 5 ZIP codes within the single-family data

3. This is an archetypical small, old, and cheap-to-build house. The 10th percentile of property size in my data is 977 sq ft; average property sizes in these ZIP codes are roughly 1200–1400 sq ft.

4. This is also why I do not use hedonic regressions to value properties, despite having the rich sales and property characteristics data to do so. Hedonic property values also yield a significant number of improvement shares above one. Using assessment values confirms that the problem is due to legacy structures and not because estimated hedonic prices are inaccurate.

with improvement shares above 0.85, the maximum being 0.89, and I assign $\tau_\omega = 1$ to all single-family properties in these ZIP codes.

2.6.3 ZIP code-level aggregate wedges

Neighborhood-level land use regulatory wedges T_{Ri} are improvement-weighted averages of plot level wedges τ_ω. As shown in the aggregation result in section 2.4,

$$T_{Ri} = \left(\frac{\sum_{\omega \in i} x_\omega / \tau_\omega}{\sum_{\omega \in i} x_\omega / \tau_\omega^{\frac{1}{1-\gamma}}} \right)^{\frac{1-\gamma}{\gamma}}.$$

Using my estimated structural values x_ω and property-level wedges τ_ω, I am able to derive ZIP-code level wedges across Cook County, as plotted in Figure 2.2.

The average T_{Ri} across Cook County ZIP codes is 0.86. Values of T_{Ri} closer to 1 denote less binding regulation. The most stringently regulated location ($T_{Ri} = 0.63$) is the affluent suburb of Evanston near Northwestern University. As shown in Figure 2.2, high T_{Ri} (low stringency) is concentrated on Chicago's South Side, where the salient issues in real estate markets are weak demand, high vacancy, and underutilized parcels rather than land-use regulations.

Downtown Chicago likewise exhibits high T_{Ri}, consistent with the prevalence of high-rise multifamily buildings and the absence of low-density zoning in the Loop. Moving north from the core, neighborhoods that combine high-value single-family homes with dense multifamily stock (e.g., the Near North Side and lakefront communities) show somewhat lower T_{Ri} (more stringency) than the CBD, and several northern suburbs have among the lowest T_{Ri} values in the county.

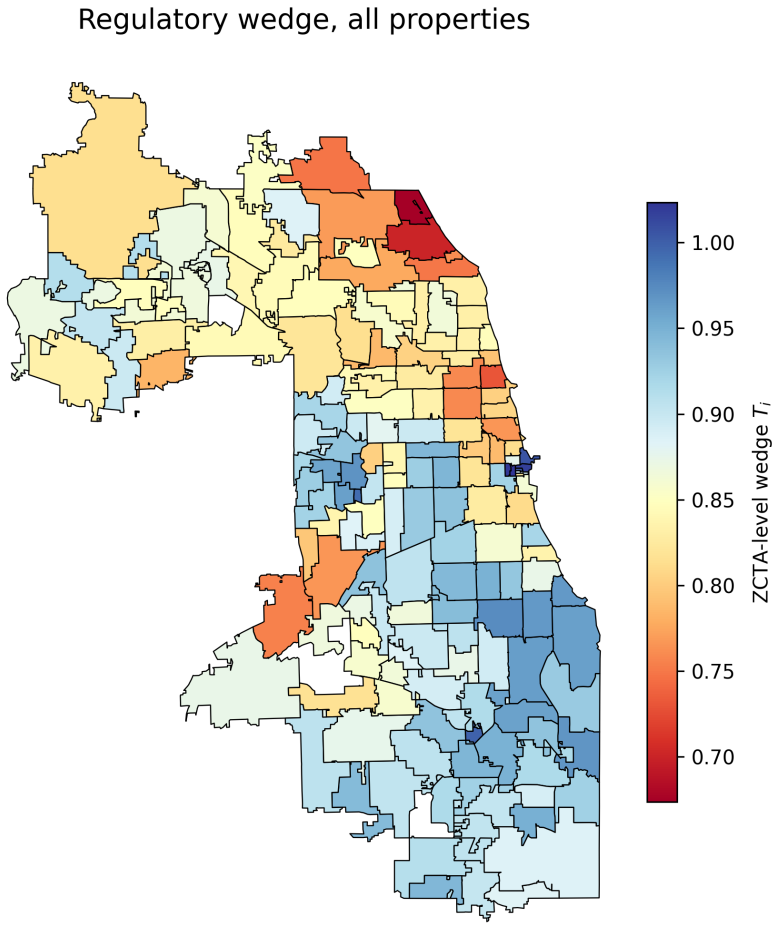

Figure 2.2: ZIP-code aggregate regulatory wedge T_{Ri} across Cook County ZIP code tabulation areas . Single family improvement shares from Davis et al. (2021), multi-family shares from RSMeans construction cost data and tax assessment values.

2.6.4 Testing wedges against zoning in Chicago

To test whether the regulatory-wedges τ_ω in my framework reflect features of land use regulation, I compare estimated [5] τ_ω for single- and multi-family properties in Chicago against their written zoning regulations. I match properties in my data to their respective zoning

5. I am able to estimate τ_ω for both single- and multi-family properties using RSMeans construction costs and assessed property values. I use Davis et al. (2021) values for single-family housing when calculating T_{Ri}. To test plot-level wedges against plot-level zoning, I use the estimated τ_ω for both single- and multi-family properties.

codes by conducting a spatial join of the property centroid coordinates to the Chicago zoning boundary map. There are 9 standard residential zoning codes in Chicago that apply to most low and medium density residential properties. Most high-density properties must undergo a review process that involves the Chicago Plan Commission and the City Council and are issued their bespoke zoning plans. I eliminate high-density properties with bespoke zoning plans as the type of land use regulations they face is not directly comparable to the 9 standard zoning codes.

Each of the 9 standard zoning codes defines a series of land use restrictions to which the property must conform. These include floor-area ratio limits, minimum lot sizes, setback requirements[6], and height limits. I examine the relationship between floor area ratio and my recovered τ_ω. Floor-area ratio (FAR) limits are one of the most studied types (Brueckner, 2007; Barr and Cohen, 2014; Cai et al., 2017; Brueckner and Singh, 2020) of land use regulations. FAR limits cap the ratio of built-up floorspace area to the size of the underlying land parcel. FAR limits represent a straightforward density and quantity limit. The 9 standard zoning codes in Chicago each come with a distinct floor-area ratio limit that ranges from 0.5 to 6.6.

Dependent variable: τ_ω	(1)	(2)	(3)	(4)
FAR ratio	0.0376^{***}	$.0396^{***}$	(Only FEs)	(Only FEs)
	(.00015)	(.00018)		
Fixed effects	None	ZIP code	Zoning code	ZIP code, Zoning code
R^2	0.107	0.189	0.124	0.217
N	482,645	482,645	482,645	482,645

*** $p < 0.01$. τ_ω is calculated using improvement shares estimated with RSMeans construction cost data. Columns (3) and (4) contain only fixed effects. The sample includes only residential properties in Chicago with the 9 standard residential zoning codes.

Table 2.2: Regression of τ_ω on FAR and zoning codes fixed effects

Columns (1) and (2) of Table 2.2 report the results of a regression of my recovered land

6. A setback requirement dictates a minimum distance between the building and the property line.

use wedges against FAR limits, with and without ZIP code fixed effects. A large floor area ratio, meaning a more lenient density limit, entails a higher τ_ω as predicted by my model. Given that regulations are not continuous in Chicago and there are only 9 categories in the standard zoning codes, I regress τ_ω on zoning code fixed effects in columns (3) and (4) of Table 2.2. I do not include FAR limits as a regressor in these regressions as each zoning code maps to one FAR; FAR limits are perfectly linear with zoning code fixed effects.

Floor area ratios and zoning-code fixed effects can account for a significant amount of the variation in τ_ω, which indicates that the price wedge is actually measuring regulations. Including zoning code fixed effects significantly raises the explanatory power of the regression. This is because properties, even when they are regulated to the same extent physically (same zoning code), face differing levels of regulatory stringency when real estate prices are different.

CHAPTER 3

A MODEL OF REGULATIONS ACROSS THE CITY

After quantifying land use regulations in Chapter 2, I model their economic effects in this chapter. It is important to consider the economic effects of land use regulations in general equilibrium; by changing local housing supply, land use regulations change prices and wages and thus alter where people live and work.

I develop a general-equilibrium quantitative spatial model (Redding and Rossi-Hansberg, 2017; Redding, 2023, 2024a) of the city in which the distribution of residents, workers, commercial activity, commuting patterns, and real estate outcomes under land use regulation are jointly determined. The model also allows me to consider and solve for *optimal* land use regulations and compare them against the status quo in Chapter 4.

In the model, I embed the housing plot–level problem of developers facing regulation into a quantitative urban general-equilibrium model. Workers choose where to live and where to work, trading off wages, housing costs, amenities, and commuting costs, while firms hire labor and produce output under perfect competition. On the supply side, price-taking housing developers face regulations in the form of price wedges and supply housing in each neighborhood, consistent with the neighborhood housing-supply relationship microfounded in Chapter 2 (see equation (2.11)). The model also allows for urban externalities through productivity and amenity spillovers: the concentration of workers and residents in a location can raise or lower local productivity and amenities.

I calibrate the model to Cook County using data on commuting, wages, and real estate from multiple sources. Inverting the model, I recover neighborhood amenities, productivities, and commuting elasticities for Cook County.

3.1 Quantitative Urban Model

The city consists of I neighborhood locations, indexed by i or j. In my application, each location is a ZIP code tabulation area. Each location i has fixed endowments of residential and commercial land, (L_{Ri}, L_{Ci}); I do not examine the Rossi-Hansberg (2004); Allen et al. (2016) policy margin of adjusting the stock of residential and commercial land within the same neighborhood. Atomistic, price-taking developers at i produce residential and commercial real estate, (H_{Ri}, H_{Ci}), taking the corresponding regulatory wedges (T_{Ri}, T_{Ci}) as given.

The city is populated by a fixed mass of $\bar{N}$ individuals. Each individual ω chooses a residence i and a workplace j. [1]. Firms at each location produce a numeraire consumption good using labor and commercial real estate. Factor markets are competitive and each factor of production is paid its marginal product. Each individual owns an equal share of the aggregate stock of residential and commercial land and thus receives land rent income. An individual's income is spent on the numeraire consumption good and on residential housing services.

3.1.1 Individual Utility

An individual chooses a residence location i and a workplace j. The individual supplies labor inelastically at their workplace j to receive wages w_j. Using these wages and their share of total land rents, the individual purchases consumption goods and residential housing at residence i. Individuals have independent and identically distributed idiosyncratic preferences $\epsilon_{ij}(\omega)$ for each residence–workplace pair $(i, j) \in I^2$ as in Monte et al. (2018).

1. In my setting, every location can serve both as a residence and a workplace; ZIP code tabulation areas are large enough geographic units such that there are no exclusively residential or commercial locations. However, the model can easily accommodate such locations.

Each location has a different level of residential amenities $B_i(\bar{B}_i, R_i) = \bar{B}_i R_i^{\phi_{amen}}$ that depends on the fundamental level of amenity $\bar{B}_i$ and the number of residents living in the location, R_i. ϕ_{amen} represents the strength of residential amenity externalities which scale with residential density. ϕ_{amen} might be positive as increased density supports more extensive local public-goods provision, richer retail variety, and stronger social capital. ϕ_{amen} might be negative due to mechanisms such as overcrowding, congestion, and increased crime due to higher residential density. [2]

Individuals incur a utility cost to commuting from residence i to workplace j, d_{ij}. The preferences of an individual ω who lives in i, works in j, consumes $C_{ij}(\omega)$ of the numeraire good and $H_{R,ij}(\omega)$ of housing is given by

$$U_{ij}(\omega) = \frac{B_i(\cdot)}{d_{ij}} \left(\frac{C_{ij}(\omega)}{\alpha} \right)^{\alpha} \left(\frac{H_{R,ij}(\omega)}{1-\alpha} \right)^{1-\alpha} \varepsilon_{ij}(\omega)\,. \tag{3.1}$$

The idiosyncratic utility shifter for each residence-workplace pair, $\varepsilon_{ij}(\omega)$, follows a Fréchet distribution with shape parameter ϵ that governs the dispersion of idiosyncratic utility. ε_{ij} follows the distribution $F(\varepsilon) = e^{-\varepsilon^{\epsilon}}$.

Individuals have two sources of income for the purchase of goods and housing: their wages w_j and the land income from their equal ownership of a city-wide land portfolio $\bar{y}$. I denote the total income of an individual working at j as $e_j = w_j + \bar{y}$. The indirect utility $V_{ij}(\omega)$ of an individual who lives at i and works at j, given residential housing prices p_{Hi}, is

$$V_{ij}(\omega) = \frac{B_i e_j}{d_{ij} p_{Hi}^{1-\alpha}} \varepsilon_{ij}(\omega)\,. \tag{3.2}$$

2. Evidence points both ways on the sign of amenity spillovers. Ahlfeldt et al. (2015); Owens et al. (2020) estimate positive reduced-form amenity externalities. Other papers find negative spillovers along specific channels, including crime (Glaeser and Sacerdote, 1999) and traffic congestion (Duranton and Puga, 2023).

3.1.2 Residence-workplace choice probabilities

Agents are ex-ante identical and differ only in their idiosyncratic preference shocks $\epsilon_{ij}(\omega)$. The probability of an individual choosing to live at i and work at j, π_{ij}, is the probability that the specific residence-workplace pair (i, j) provides the individual with the highest level of utility across all possible residence-workplace combinations, given the realization of the individual's idiosyncratic preference shocks $\{\epsilon_{ij}(\omega)\}_{(i,j)\in I^2}$.

$$\pi_{ij} = P\left(V_{ij} > \max_{(m,n)\neq(i,j)} V_{mn}\right)$$

Using standard manipulations of the Fréchet distribution (Eaton and Kortum, 2002; Redding, 2023), the residence-workplace pair choice probabilities π_{ij} are given by

$$\pi_{ij} = \frac{\left(B_i e_j/(d_{ij} p_{Hi}^{1-\alpha})\right)^{\epsilon}}{\sum_{m,n\in I^2}\left(B_m e_n/(d_{mn} p_{Hm}^{1-\alpha})\right)^{\epsilon}}\,. \tag{3.3}$$

Ex-ante expected utility $\bar{U}$ across all locations, given shape parameter ϵ of the distribution of Fréchet idiosyncratic shocks $\varepsilon_{ij}(\omega)$, is

$$\bar{U} = \Gamma\left(\frac{\epsilon-1}{\epsilon}\right)\left(\sum_{i,j}\left(B_i e_j/(d_{ij} p_{Hi}^{1-\alpha})\right)^{\epsilon}\right)^{1/\epsilon} \tag{3.4}$$

, where $\Gamma(\cdot)$ is the gamma function.

3.1.3 Production

In each location i, there is a continuum of identical, atomistic, perfectly competitive firms that produce a freely traded numeraire good using labor and commercial real estate with a constant returns to scale Cobb-Douglas production function. The continuum of firms can

be represented by a single representative firm in each location. The representative firm in location i chooses labor N_i and commercial real estate H_{Ci} to maximize profits, given the local wage w_i and the local commercial real estate price p_{Ci}.

Output in location i is given by a constant-returns-to-scale Cobb–Douglas production function:

$$Y_i \;=\; A_i(\bar{A}_i, N_i)\, N_i^{\beta}\, H_{Ci}^{1-\beta}, \qquad A_i(\bar{A}_i, N_i) \;=\; \bar{A}_i\, N_i^{\phi_{\text{prod}}}. \tag{3.5}$$

$\bar{A}_i$ is exogenous baseline productivity in location i, and ϕ_{prod} captures a productivity externality: the concentration of workers N_i in location i raises productivity A_i of all individual firms in the location through agglomeration forces such as knowledge spillovers and the development of business infrastructure. Each firm is atomistic and therefore does not internalize how its own hiring affects overall productivity, so $A_i(\cdot)$ is taken as given in the representative firm's input choice.

Profit maximization implies that labor is paid its marginal product, so wages in location i are

$$w_i \;=\; A_i(\cdot)\,\beta \left(\frac{H_{Ci}}{N_i}\right)^{1-\beta}. \tag{3.6}$$

Similarly, the inverse demand for commercial real estate in location i is

$$p_{Ci} \;=\; A_i(\cdot)\,(1-\beta)\, \left(\frac{N_i}{H_{Ci}}\right)^{\beta}. \tag{3.7}$$

3.1.4 Real Estate

In each location i, there is an exogenous endowment of residential and commercial land quantities, L_{Ri} and L_{Ci}. Price-taking atomistic landlords use the numeraire good and land to produce residential and commercial buildings, given rents p_{Ri} and p_{Ci}.

There exist land use regulations τ_ω on individual land parcels that decrease the marginal

revenue product of real estate, and thus decrease real estate supply. Given the aggregation result in Section 2.4, neighborhood level residential and commercial real-estate supply H^S_{Ri} and H^S_{Ci} can be expressed in terms of aggregate, neighborhood level wedges T_{Ri}, T_{Ci} as

$$H^S_{Ri}(p_{Ri}) = [\gamma_R T_{Ri}]^{\frac{\gamma_R}{1-\gamma_R}} Z_{Ri}^{\frac{1}{1-\gamma_R}} L_{Ri} \cdot p_{Ri}^{\frac{\gamma_R}{1-\gamma_R}}\,, \tag{3.8}$$

$$H^S_{Ci}(p_{Ci}) = (\gamma_C T_{Ci})^{\frac{\gamma_C}{1-\gamma_C}} Z_{Ci}^{\frac{1}{1-\gamma_C}} L_{Ci} \cdot p_{Ci}^{\frac{\gamma_C}{1-\gamma_C}}\,. \tag{3.9}$$

Given the Cobb-Douglas individual utility over the numeraire and real estate, total expenditure on real estate at i is fraction $(1-\alpha)$ of total income of all residents in that location. Demand for residential real estate at i, H^D_{Ri} at location i is thus

$$H^D_{Ri}(p_{Ri}) = \frac{(1-\alpha)\bar{N}\sum_j \pi_{ij}(w_j + \bar{y})}{p_{Ri}} = \frac{(1-\alpha)\bar{v_i}R_i}{p_{Ri}}\,. \tag{3.10}$$

where $\bar{v_i}$ is the average total income across all residents living in location i. Given Cobb-Douglas production by price-taking producers, total commercial real estate revenue at i is fraction $(1-\beta)$ of producer output. Thus, demand for commercial real estate, H^D_{Cj}, is given by

$$H^D_{Cj}(p_{Cj}) = \frac{1-\beta}{\beta}\frac{w_j N_j}{p_{Cj}}\,. \tag{3.11}$$

The fraction of total residential real estate revenue paid to improvements is $\gamma_R T_{Ri}$, so landlord income from residential real estate equals $(1-\gamma_R T_{Ri})\, p_{Ri} H_{Ri}$. In the commercial market, the corresponding improvement share and landlord income are $\gamma_C T_{Ci}$ and $(1-\gamma_C T_{Ci})\, p_{Ci} H_{Ci}$. As discussed previously, tighter land use regulation (lower T_{Ri} or T_{Ci}) shifts revenue from improvements to land, potentially raising land incomes.

Each individual, regardless of where they live or work, owns an equal share of all residential and commercial land in the city. They receive land income $\bar{y}$ in addition to their work

wage, given by

$$\bar{y} = \frac{1}{\bar{N}} \left[\sum_i (1 - \gamma_C T_{Ci}) \cdot p_{Ci} H_{Ci} + \sum_i (1 - \gamma_R T_{Ri}) \cdot p_{Ri} H_{Ri} \right] . \tag{3.12}$$

3.1.5 Adding up conditions

To complete the equilibrium, we need two simple accounting identities to make sure that the average incomes and number of workers and residents conform with residence-workplace choice probabilities π_{ij}. Average income $\bar{v_i}$ of residents is given by:

$$\bar{v_i} = \frac{\sum_j \pi_{ij}(w_j + \bar{y})}{\sum_j \pi_{ij}} . \tag{3.13}$$

The residential population R_i at each location must be consistent with choice probabilities

$$R_i = \sum_j \pi_{ij} \bar{N} . \tag{3.14}$$

Worker population at each location must also be consistent with choice probabilities

$$N_j = \sum_i \pi_{ij} \bar{N} . \tag{3.15}$$

3.1.6 Equilibrium

Given fundamentals $\{\bar{A}_i, \bar{B}_i, Z_{Ri}, Z_{Ci}, L_{Ri}, L_{Ci}, T_{Ri}, T_{Ci}\}_{i \in I}$, commuting costs $\{d_{ij}\}_{i,j \in I^2}$, parameters $\phi_{\text{amen}}, \phi_{\text{prod}}, \alpha, \beta, \gamma_R, \gamma_C, \epsilon$, and $\bar{N}$, a spatial general equilibrium is a set of endogenous objects $\{\pi_{ij}, w_i, p_{Ri}, p_{Ci}, H_{Ri}, H_{Ci}, R_i, N_i, \bar{v}_i, \}_{i \in I}$ and $\bar{y}$ such that:

(i) **Individual optimization:** Individuals choose a residence i and workplace j to maxi-

mize utility $V_{ij}(\omega)$ given wages w_j, housing prices p_{Ri}, amenities B_i, commuting costs d_{ij}, and land income $\bar{y}$. This implies the residence-workplace choice probabilities given in (3.3).

(ii) **Firm optimization:** In each i, the representative firm chooses labor N_i and commercial real estate H_{Ci} to maximize profits under (3.5). First-order conditions determine wages w_i and commercial real-estate prices p_{Ci} as in (3.6) and (3.7).

(iii) **Real estate market clearing:** In each location, the demand and supply for residential and commercial real estate equate. The supply of residential and commercial real estate is given by (3.8) and (3.9). The demand for residential and commercial real estate is given by (3.10) and (3.11).

(iv) **Adding up.** Residential and worker populations, along with total incomes from wages and land rents, are consistent with π_{ij}: R_i and N_i satisfy (3.14) and (3.15); $\bar{v}_i$ satisfies (3.13); $\bar{y}$ satisfies (3.12).

The equilibrium can be written in a more concise system of $5 \times I$ equations in the endogenous variables $X = \{p_{Ri}, p_{Ci}, w_i, R_i, N_i\}_{i \in I}$. The 5 equations that describe the general equilibrium for each location i are given by

$$[\gamma_R T_{Ri}]^{\frac{\gamma_R}{1-\gamma_R}} Z_{Ri}^{\frac{1}{1-\gamma_R}} L_{Ri}\, p_{Ri}^{\frac{\gamma_R}{1-\gamma_R}} \;=\; \frac{1-\alpha}{p_{Ri}}\, \bar{N} \sum_j \pi_{ij}(R_i, p_{Ri}) \left(w_j + \bar{y}\right) \qquad \text{(Res. RE clearing)} \tag{3.16}$$

$$(\gamma_C T_{Ci})^{\frac{\gamma_C}{1-\gamma_C}} Z_{Ci}^{\frac{1}{1-\gamma_C}} L_{Ci}\, p_{Ci}^{\frac{\gamma_C}{1-\gamma_C}} \;=\; \frac{1-\beta}{\beta}\, \frac{w_i N_i}{p_{Ci}} \qquad \text{(Com. RE clearing)} \tag{3.17}$$

$$w_i \;=\; \bar{A}_i\, N_i^{\phi_{\text{prod}}}\, \beta \left(\frac{1-\beta}{\beta} \cdot \frac{w_i}{p_{Ci}} \right)^{1-\beta} \qquad \text{(Labor mkt. clearing)} \tag{3.18}$$

$$R_i \;=\; \bar{N} \sum_j \pi_{ij}(R_i, p_{Ri}) \qquad \text{(Residents adding up)} \tag{3.19}$$

$$N_i \;=\; \bar{N} \sum_j \pi_{ji}(R_j, p_{Rj}) \qquad \text{(Workers adding up)} \tag{3.20}$$

The residence-workplace choice probability $\pi_{ji}(R_j, p_{Rj})$ is expressed in (3.19) and (3.20) in terms of the endogenous variables R_i and p_{Ri} as

$$\pi_{ij}(R_i, p_{Ri}) = \frac{\left(\bar{B}_i R_i^{\phi_{\text{amen}}}\, e_j / \left(d_{ij}\, p_{Ri}^{1-\alpha} \right) \right)^{\epsilon}}{\sum_{m,n \in I^2} \left(\bar{B}_m R_m^{\phi_{\text{amen}}}\, e_n / \left(d_{mn}\, p_{Rm}^{1-\alpha} \right) \right)^{\epsilon}}$$

, where $e_j = w_j + \frac{1}{\bar{N}} \left[\sum_i (1 - \gamma_C T_{Ci})\, p_{Ci} H_{Ci} + \sum_i (1 - \gamma_R T_{Ri})\, p_{Ri} H_{Ri} \right]$,.

The five equations (3.16), (3.17), (3.18), (3.19), and (3.20) hold for every location i in equilibrium. Stacking these across all $i \in I$ yields a $5 \times I$ nonlinear system, which we can write as $F(X, T_R) = 0$ for a given vector of residential land use wedges $T_R = \{T_{Ri}\}_{i \in I}$ that

all agents take as given.

3.2 Model Calibration

To calibrate the model in order to conduct counterfactual experiments, I recover the model fundamentals of regulatory wedges, location amenities, final goods productivity, and real estate productivity $\{T_{Ri}, \bar{A}_i, \bar{B}_i, Z_{Ri}, Z_{Ci}\}$ for ZIP code tabulation areas (ZCTAs) in Cook County, Illinois for the year 2019. Cook County contains the City of Chicago and surrounding municipalities and suburbs (e.g., Evanston, Oak Park, Schaumburg, Arlington Heights, Elgin). These jurisdictions form a tightly integrated labor and land market, with substantial commuting across locations within the County.

Parameter	**Value**	**Source**
α	0.7	Davis and Ortalo-Magné (2011)
β	0.80	Ahlfeldt et al. (2015)
γ_C, γ_R	0.85	Ahlfeldt and McMillen (2014); Epple et al. (2010)
ϵ	4.31	Commuting gravity regression
ϕ_{prod}	0.07	Ciccone and Hall (1996); Ahlfeldt et al. (2015)
ϕ_{amen}	0.15, -0.04, -0.15	Ahlfeldt et al. (2015); Glaeser and Sacerdote (1999); Duranton and Puga (2023)
$\bar{N}$	1,730,851	Total full-time workers in Cook County, LODES

Table 3.1: Parameter Values and Sources

Table 3.1 reports the parameters that I use in my calibration. The model parameters $\{\alpha, \beta, \gamma_C, \gamma_R\}$ are matched to estimates from the literature, implying a 30% real-estate expenditure share of total income, an 80% wage share of output, and an improvement share of real estate value of 85%. I use the estimate of the strength of productivity spillovers of $\phi_{prod} = 0.07$ from Ciccone and Hall (1996) and Ahlfeldt et al. (2015). While the literature has agreed on positive productivity spillovers, there has been empirical evidence for both positive and negative amenity spillovers. I explore policies under both positive and negative externality spillovers.

Data from the LEHD Origin-Destination Employment Statistics (LODES), ZIP-code business patterns, and travel times allow me to observe the endogenous outcomes in the model, such as wages and commuting. I then use these observed outcomes to recover fundamentals $\{\bar{A}_i, \bar{B}_i, Z_{Ri}, Z_{Ci}\}$ via the process of "model inversion" (Redding and Rossi-Hansberg, 2017). I use the neighborhood-level residential regulatory price wedges T_{Ri} for Cook County ZCTAs measured in Chapter 2.6, which are derived from CoreLogic, RSMeans, and Davis et al. (2021) data.

I quantify the model at the ZIP Code Tabulation Area (ZCTA) level; each location i and j in my calibration refers to a ZCTA. ZCTAs are statistical geographic units created by the U.S. Census Bureau to approximate United States Postal Service (USPS) ZIP codes. ZIP codes are defined by the USPS for mail delivery and can be non-contiguous, overlapping, and subject to frequent boundary changes. ZCTAs are constructed by the Census Bureau to provide a stable, census-compatible geography that approximates ZIP codes.

ZCTAs do not align perfectly with county boundaries. I restrict my sample to ZCTAs whose land area lies mostly within Cook County and drop ZCTAs that extend into Cook County but are primarily in adjacent counties (Lake, DuPage, or Kane), yielding a final sample of $I = 175$ ZCTAs. At the ZCTA-level geographical resolution, there are no purely residential or commercial locations. Every ZCTA has positive amounts of workers and residents and both residential and commercial real estate. Each ZCTA serves as both a residence and workplace.

3.3 Residents, workers, wages, commuting

I estimate commuting elasticities and location fundamentals using commuting and income data. I obtain origin–destination flows π_{ij}, residents R_i, and workplace employment N_i from the Longitudinal Employer–Household Dynamics Origin–Destination Employment Statistics

(LODES). I construct census block-level π_{ij}, R_i, and N_i directly from the census block-to-census block residence–workplace flows in LODES.

ZCTAs are constructed such that each census block maps to exactly one ZCTA. I use the Census Bureau ZIP Code Tabulation Areas Relationship Files to map census blocks to ZCTAs. The census block-to-ZCTA crosswalk files, along with the LODES data, allow me to construct the commuting matrix at the ZCTA level.

ZCTA-level wages, w_i, are obtained from the U.S. Census Bureau's 2019 ZIP Code Business Patterns (ZBP) dataset. The ZBP dataset contains the number of establishments, employment, and annual payrolls at the ZIP code level. I obtain ZIP-code level annual wages by dividing aggregate payrolls by employment. The establishment payroll data in the ZBP dataset allows me to confirm that I am indeed obtaining worker wages, rather than resident incomes; commuting in my model requires a careful distinction between these two concepts.

I measure bilateral commuting costs d_{ij} as driving times between ZCTAs i and j. Because ZCTAs are relatively coarse units, I aggregate data from census-block pairs: for each ZCTA pair, I obtain a commuting-flow-weighted average of block-to-block shortest-path driving times (during 9 a.m. on a Wednesday) using the Conveyal R5 routing engine and OpenStreetMap road network data. Since census blocks are nested within ZCTAs by construction, this yields

$$d_{ij} = \sum_{b \in i} \sum_{c \in j} \omega_{bc}\, t_{bc}\,,$$

where t_{bc} is the driving time from block b to block c, and ω_{bc} are weights proportional to commutes from b to c (normalized to sum to one for each (i, j)). I set $t_{bb} = 1$ minute for trips that begin and end in the same block. Bilateral commuting costs are not symmetrical; $d_{ij} \neq d_{ji}$.

3.4 Estimating commuting elasticity

I estimate ϵ, the elasticity of commuting flows with respect to commuting costs d_{ij}, using the observed commuting flows for Cook County in LODES. Taking logs of the commuting probability equation (3.3) yields the standard gravity equation in commuting,

$$\log(\pi_{ij}) = -\epsilon \log(d_{ij}) + \epsilon \log\left(\frac{B_i}{p_{Hi}^{1-\alpha}}\right) + \epsilon \log(e_j) + log(\Phi)$$

,where $\Phi = \sum_{m,n \in I^2} \left(B_m e_n / (d_{mn} p_{Hm}^{1-\alpha})\right)^{\epsilon}$ is the denominator of the commuting probability equation. I recover ϵ using the estimating equation

$$\log(\pi_{ij}) = -\epsilon \log(d_{ij}) + \kappa_i + \nu_j + \varsigma_{ij} \, .$$

κ_i is an origin i fixed effect and ν_j is a destination j fixed effect. An OLS estimate of ϵ is generally inconsistent and biased due to the elimination of zero-flow observations from the regression, as highlighted in Silva and Tenreyro (2006). Following their approach to estimating gravity equations, I use the Poisson Pseudo–Maximum Likelihood(PPML) estimator, which uses information on zero flows and allows for heteroskedasticity in the error term ς_{ij} .

Table 3.2: Commuting elasticity estimates

	OLS	PPML
$\log(\pi_{ij})$	-3.01	-4.31
	(0.010)	(0.012)
Fixed effects	home, workplace	home, workplace
R^2/ PPML pseudo-R^2	0.41	0.46
N	27148	30625

Table 3.2 reports OLS and PPML estimates of ϵ. Given that ZCTAs are relatively coarse geographies, there are only 3447 zero ZCTA-to-ZCTA flows out of the $175 \times 175 = 30625$

origin-destination ZCTA pairs for Cook County. OLS is an inconsistent estimator of ϵ even without zeros due to heteroskedasticity. Assuming that $\mathbb{E}\left(e^{\varsigma_{ij}} \mid \kappa_i, \nu_j, d_{ij}\right) = 1$, PPML yields consistent estimates of ϵ. I set $\epsilon = 4.31$ using the PPML estimate. This tracks closely to estimated commuting elasticity values ranging from 4 to 6 in Ahlfeldt et al. (2015) and Monte et al. (2018).

3.5 Local neighborhood characteristics

Given observed equilibrium wages, commuting, and real estate quantities, I am able to back out local neighborhood productivities $A_i(\cdot)$, amenities $B_i(\cdot)$, and residential and commercial real estate supply parameters Z_{Ri}, Z_{Ci} by inverting the model. I obtain land areas and real estate quantities $H_{Ri}, H_{Ci}, L_{Ri}, L_{Ci}$ from the CoreLogic data. Although I can observe real estate prices p_{Ri} and p_{Ci}, using them directly would imply real estate expenditures that are inconsistent with the Cobb-Douglas utility and production functions used in the model. I therefore recover model-consistent real estate prices from wages and the equilibrium conditions. I use firms' demand function for commercial real estate (3.11) to write p_{Ci} in terms of observed workers N_i, wages w_i, and commercial real estate quantity H_{Ci}:

$$p_{Ci} = \frac{1-\beta}{\beta} \frac{w_i N_i}{H_{Ci}} \,. \tag{3.21}$$

The Cobb-Douglas expenditure shares given by equation (3.10) cannot be used to solve directly for residential real estate prices p_{Ri}. This is because consumers' income consists not only of observed wages w_j, but also of each agent's share of the city-wide land portfolio $\bar{y}$, which is itself an endogenous function of p_{Ri}. Using (3.10) to derive the inverse demand function for residential real estate at each location yields

$$p_{Ri} = \frac{(1-\alpha)\bar{N}\sum_j \pi_{ij}(w_j + \bar{y})}{H_{Ri}}\,.$$

Substituting in $\bar{y} = \frac{1}{\bar{N}}\sum_i (1-\gamma_C T_{Ci})\cdot p_{Ci}H_{Ci} + \sum_i (1-\gamma_R T_{Ri})\cdot p_{Ri}H_{Ri}$ into the inverse demand function for residential real estate, I get

$$p_{Ri} = \frac{(1-\alpha)\bar{N}\sum_j \pi_{ij} w_j}{H_{Ri}} + \frac{(1-\alpha)\sum_j \pi_{ij}}{H_{Ri}} \sum_i \left[(1-\gamma_C T_{Ci})p_{Ci}H_{Ci} + (1-\gamma_R T_{Ri})p_{Ri}H_{Ri}\right]. \tag{3.22}$$

Equation (3.22) describes a system of I simultaneous linear equations (one for each location) in the I residential real estate prices at each location $\{p_{Ri}\}_{i\in I}$. I solve the system of equations to obtain residential real estate prices at each location. The other objects in the system of equations are either observed in the data ($\pi_{ij}, H_{Ri}, H_{Ci}, w_j$), or are recovered from the observed equilibrium (p_{Ci}, using equation (3.21)).

With both quantities H_{Ri}, H_{Ci} and prices p_{Ri}, p_{Ci} of real estate, I use the real estate supply function to recover residential and commercial real estate supply productivities Z_{Ri}, Z_{Ci}. Rearranging equations (3.8) and (3.9),

$$Z_{Ri} = \frac{1}{\gamma_R T_{Ri}} \left(\frac{H_{Ri}^S}{L_{Ri}\, p_{Ri}^{\gamma_R/(1-\gamma_R)}} \right)^{\frac{1-\gamma_R}{\gamma_R}} \quad , \quad Z_{Ci} = \frac{1}{\gamma_C T_{Ci}} \left(\frac{H_{Ci}^S}{L_{Ci}\, p_{Ci}^{\gamma_C/(1-\gamma_C)}} \right)^{\frac{1-\gamma_C}{\gamma_C}}\,.$$

I now turn to recovering composite amenities and productivities $A_i = \bar{A}_i N_i^{\phi_{prod}}$ and $B_i = \bar{B}_i R_i^{\phi_{amen}}$. Fundamental productivities $\bar{A}_i$ and amenities $\bar{B}_i$ for given values of ϕ_{prod} and ϕ_{amen} can then be recovered using the number of residents R_i and workers N_i within

each ZCTA. To recover composite productivity A_j, I rearrange equation (3.5) as

$$A_j = \frac{Y_j}{N_j^{\beta} H_{Cj}^{1-\beta}}. \tag{3.23}$$

The value of final goods produced at each location can be recovered using total wages; $Y_i = \frac{N_i W_i}{\beta}$. To recover composite residential amenities B_i, I use the commuting probability equation (3.3) to express the number of residents in each location as

$$R_i = \frac{\left(\frac{B_i}{p_{Hi}^{1-\alpha}}\right)^{\epsilon} \sum_j \left(\frac{e_j}{d_{ij}}\right)^{\epsilon}}{\sum_{m,n \in I^2} \left(B_m e_n / (d_{mn} p_{Hm}^{1-\alpha})\right)^{\epsilon}}. \tag{3.24}$$

The term $\sum_j \left(\frac{e_j}{d_{ij}}\right)^{\epsilon}$ is a resident market access (RMA) term that reflects the travel cost-weighted access of residents in location i to workplaces. Rearranging the equation gives the following ratio of residential amenities between locations i and k,

$$\frac{B_i}{B_k} = \frac{R_i^{1/\epsilon} p_{Ri}^{1-\alpha} \left(\sum_j \left(\frac{e_j}{\tau_{kj}}\right)^{\epsilon}\right)^{1/\epsilon}}{R_k^{1/\epsilon} p_{Rk}^{1-\alpha} \left(\sum_j \left(\frac{e_j}{d_{ij}}\right)^{\epsilon}\right)^{1/\epsilon}}. \tag{3.25}$$

This allows identification of amenities up to scale. I set amenities for the first location in my dataset to 1. The choice of normalization does not matter since amenities are utility shifters.

Figure 3.1 plots the distribution of composite productivity and amenities across Cook County. The relative distribution of fundamental productivities and amenities $\bar{A}_i, \bar{B}_i$ does not change significantly for the values of ϕ_{amen} and ϕ_{prod} that I use in my calibration.

Downtown Chicago, O'Hare international airport, and the affluent northern and western suburbs of Cook County contain many of the high productivity locations within the county.

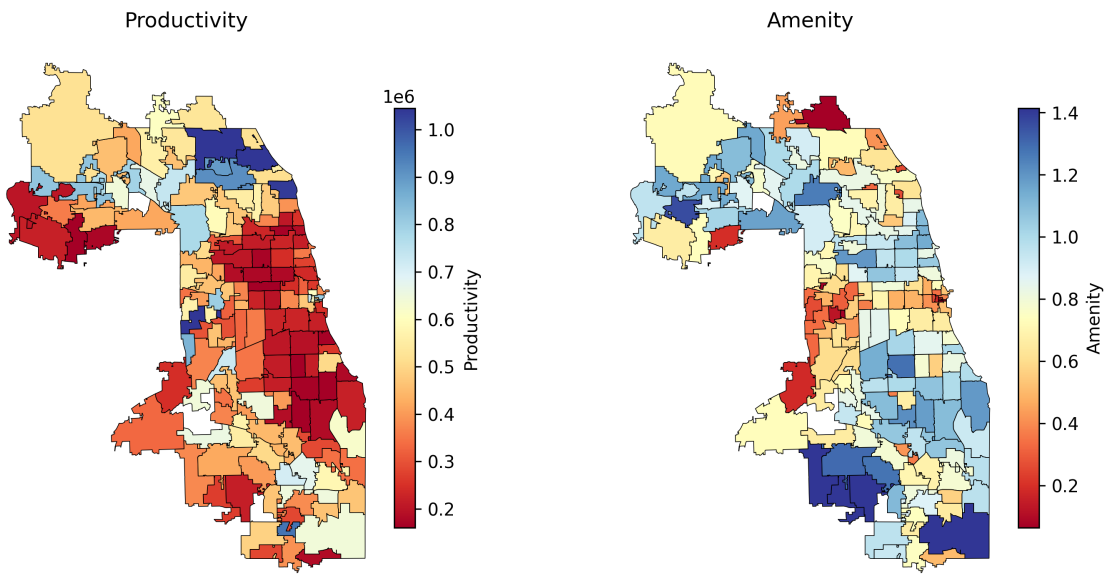

Figure 3.1: Composite productivities and amenities across ZCTAs in Cook County. Productivities are on the 10^6 scale as I denominate real estate in millions of square feet. Gaps in the map are forest preserves not included in ZCTA shapefiles.

The most productive ZIP code in my sample is ZIP code 60141, which consists of the campus of Loyola University Medical Center and a small amount of residential real estate. ZIP Code Business Patterns data report an average wage of $125,668 there, far above the Cook County average of $69,369. Given the high salaries in medical professions, the productivity of this hospital-dominated ZIP code is not surprising. Outside downtown Chicago, very productive areas are generally located in the northern and western suburbs: Evanston, Glencoe, and Winnetka in the northern suburbs, and Western Springs and Westchester in the west.

There is a negative correlation between productivity and amenities of -0.32; locations with high productivities tend to have lower levels of amenities. This is because in spatial equilibrium, expected utilities must equate across residence locations. The distribution of productivities is generally spatially correlated; locations with high productivity tend to be have high-productivity neighbors. As such, high productivity locations tend to have higher residential market access(RMA) as they are nearer to other high-wage, high productivity

locations. Absent a large number of residents due to high residential market access, model inversion would impute that the location would have low amenities to rationalize the residential population given high RMA.

The distribution of amenities might not totally conform to general perceptions of residential amenities in Cook County. The model inversion yields high amenities in the South and Southwest sides of Chicago, which are predominantly African American and Hispanic neighborhoods with lower than average wages and higher than average crime. All individuals in the model are ex-ante identical. The amenities obtained via model inversion are residuals that rationalize the population of poor South Side neighborhoods; given low access to high-wage jobs, the model infers high amenities to rationalize why individuals in the South Side have not mostly moved to the affluent suburbs or downtown. Accounting for the pattern of racial segregation in Cook County would require modifications to the model involving non-homothetic preferences and heterogeneity in race and education across individuals. I abstract from such considerations to keep the model parsimonious, given the computational difficulty of solving the optimal land use regulation problem.

To better understand the spatial linkages that shape equilibrium outcomes, I examine resident market access (RMA) and firm market access (FMA) across Cook County. RMA at location i is the commuting-cost-weighted access of residents to income-earning workplaces, while FMA at location j is the commuting-cost-weighted access of firms to worker-attractive residential locations. These objects summarize two distinct spatial forces in the model: the ease with which residents can access high-paying jobs, and the ease with which firms can access the residential locations from which they draw labor. As derived in Appendix A.4, the expressions for RMA and FMA are

$$RMA_i \equiv \sum_j \left(\frac{e_j}{d_{ij}} \right)^{\epsilon}, \qquad FMA_j \equiv \sum_i \left(\frac{B_i}{d_{ij} p_{Ri}^{1-\alpha}} \right)^{\epsilon}. \tag{3.26}$$

Figure 3.2 plots the spatial distribution of RMA and FMA across Cook County. Downtown Chicago is the location with the highest levels of both resident and firm market access, reflecting its central position in the county's transportation network and its strong connectivity to the rest of the metropolitan area. In this sense, downtown is the most spatially connected part of Cook County: it is both the easiest place from which residents can access jobs and one of the easiest places for firms to draw workers from across neighborhoods.

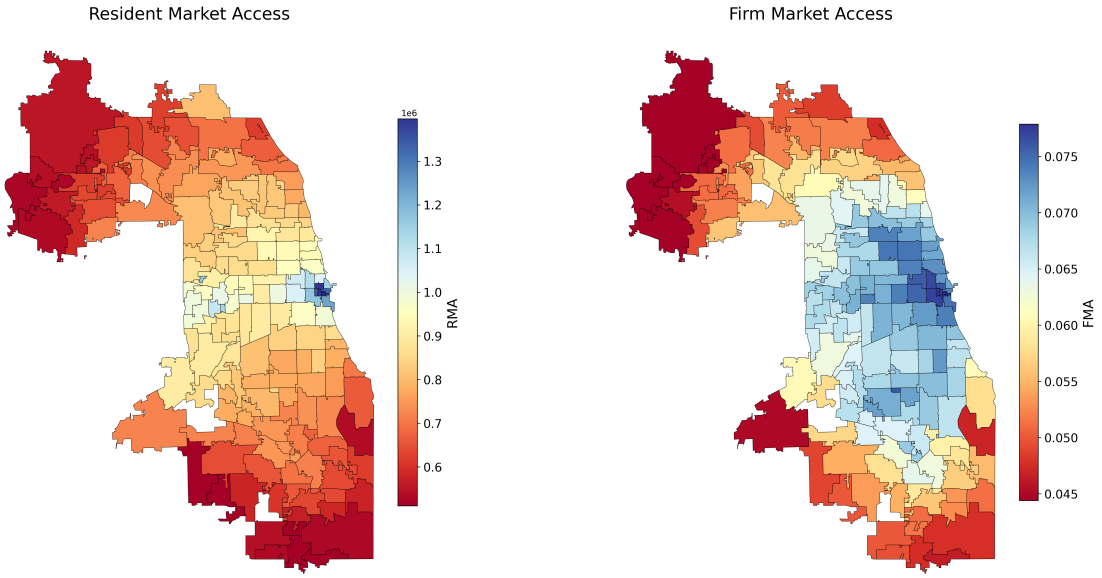

Figure 3.2: Resident market access (left) and firm market access (right) across ZIP codes in Cook County.

The spatial gradients of RMA and FMA away from downtown are, however, quite different. RMA declines sharply with distance from downtown. This is because RMA measures access to jobs, and high-paying jobs are disproportionately concentrated in downtown Chicago and a small number of nearby employment centers. As a result, neighborhoods farther from the urban core experience a rapid fall in residential access to earnings opportunities. FMA, by contrast, declines much more gradually. FMA measures firms' access to workers, and the residential population of Cook County is much more evenly spread across neighborhoods than employment is. Firms therefore retain access to a broad labor pool over

a much wider geographic area, so firm market access tapers off more slowly than resident market access.

These differences are important for understanding the geography of optimal land use regulations. Locations with high RMA are attractive residential locations because they provide strong access to income, while locations with high FMA are attractive workplace locations because they provide firms with access to labor. Optimal regulation depends not only on local productivity and amenities, but also on these broader spatial linkages across neighborhoods. I examine this link in Chapter 4.

CHAPTER 4

COMPUTING OPTIMAL LAND USE POLICIES

Given the parameter values and location fundamentals for Chicago that I have recovered in chapter 3, I can compute counterfactual equilibrium outcomes at any desired regulation vector $\{T_{Ri}\}_{i \in I}$. I then optimize over $\{T_{Ri}\}$ to find the government's optimal residential land use policy.

Existing evaluations of land use regulation typically benchmark current rules against *predefined* policy counterfactuals rather than *optimal*, welfare-maximizing policy. Common exercises relax regulations to a historical baseline or remove them wholesale (Hsieh and Moretti, 2019; Herkenhoff et al., 2018; Duranton and Puga, 2023; Babalievsky et al., 2023; Mei, 2022; Macek, 2024). The empirical setting for these existing works is mainly the New York or San Francisco area; given the stringency of regulation in these cities, almost any type of deregulation yields sizable gains compared to the status quo. These counterfactuals are not designed to pin down the welfare-maximizing spatial pattern of regulation across locations.

Urban externalities in amenities and production, as modeled in chapter 3, mean that a laissez-faire policy of complete regulation is not optimal. Recent work on optimal, welfare-maximizing spatial policies (Rossi-Hansberg et al., 2025; Fajgelbaum and Gaubert, 2020) emphasizes correcting amenity and productivity externalities with location-based taxes and subsidies. These location-specific transfers may be practically and politically feasible. In contrast, land use regulations are commonly adopted by local governments to address urban externalities. As such, I consider what happens when governments can only impose land use regulations to correct urban externalities.

4.1 Solving for equilibrium for a given level of land use regulations

Given the parameter values and location fundamentals from the previous section, I can compute counterfactual equilibrium outcomes at any given level of regulations by solving for economic variables X that solve the general equilibrium system of equations $F(X, T_R)$. In more formal terms, given a distribution of land use regulations $\{T_{Ri}\}_{i \in I}$, an equilibrium is a vector of endogenous variables $X = \{p_{Ri}, p_{Ci}, w_i, R_i, N_i\}_{i \in I}$ that solves the system $F(X, \{T_{Ri}\}) = 0$. The general-equilibrium system $F(X, \{T_{Ri}\})$ is a system of five equations (3.17), (3.16), (3.18), (3.19), (3.20) for each location i, for a total of $5I = 5 \times 175$ equations in $5I$ unknowns.

4.2 Optimal land use regulation problem

I define the problem of a local government which sets residential land use regulations $T_R = \{T_{Ri}\}_{i \in I}$ to maximize aggregate welfare $O(\cdot)$. In my setting, I impose that the policymaker maximizes aggregate welfare across all agents, which is also the expected ex-ante utility across all agents. My solution framework is robust to other specifications of the government's objective function. Using equation (3.4) , the policymaker's welfare objective $O(X, T_R)$ is

$$O(X, T_R) \;=\; \left(\sum_{i,j} \left(\frac{B_i \, e_j}{d_{ij} \, p_{Hi}^{1-\alpha}} \right)^{\epsilon} \right)^{1/\epsilon}.$$

The planner chooses land use regulations across neighborhoods T_R; each regulatory policy T_R entails a different realization of endogenous economic outcomes X that satisfy $F(X, T_R) = 0$.

The policymaker's optimal land use regulation problem is

$$\begin{aligned} \max_{T_R, X} \quad & O(X, T_R) \\ \text{s.t.} \quad & F(X, T_R) = 0, \\ & 0 \leq T_{Ri} \leq 1 \quad \forall i. \end{aligned} \tag{4.1}$$

The choice variables in this optimal policy problem include both T_R and X. For any chosen T_R, the endogenous variables X must be consistent with general equilibrium $F(X, T_R) = 0$. The bounds on T_{Ri} reflect (i) a complete ban on building at $T_{Ri} = 0$ and (ii) a minimal-regulation case at $T_{Ri} = 1$ in which developers face no land use regulations .[1] I ignore the numerical bounds on T_{Ri} in my theoretical exposition for the purposes of simplifying notation.

4.2.1 Land use regulations are not first-best policy for correcting urban externalities, and admit no analytical solution

Fajgelbaum and Gaubert (2020) show that in the general class of quantitative urban models with urban externalities, of which my model is a member, the no-intervention equilibrium is inefficient because individuals do not internalize how residence–workplace choices affect others via amenity and productivity spillovers. Fajgelbaum and Gaubert (2025) show that in constant-spillover-elasticity spatial models with commuting (encompassing my setting), the first best can be implemented by ad valorem labor subsidies and residence-specific lump-sum taxes. While theoretically optimal, such people-based instruments are often not politically feasible.

1. Values $T_{Ri} > 1$ would correspond to location-wide construction subsidies. Such policies may exist for specific projects but not at the ZIP/neighborhood scale I study, so I preclude $T_{Ri} > 1$ in my second-best problem.

A common justification for land use regulation is to address urban externalities. I show that land use regulations are *not* first-best instruments for that purpose.

Proposition 1. Suppose the government can set residential and commercial land use regulations $\{T_{Ri}, T_{Ci}\}$ at each location. In the presence of urban externalities ($\phi_{\text{amen}}, \phi_{\text{prod}} \neq 0$), $\{T_{Ri}\}$ and $\{T_{Ci}\}$ cannot attain the first-best, aggregate welfare-maximizing allocation because they cannot *perfectly* internalize amenity and productivity externalities.

Proof. See Appendix A.1. I consider the first-best planner's allocation problem (choosing consumption and housing directly subject to resource constraints) and compare it to the decentralized equilibrium with regulation. I show that the two cannot coincide. □

Consider the case of residential real estate regulations. The crux of the proof is that in a first-best world, housing should be built until the social cost of an extra unit of improvements equals the social benefit of using an extra unit of improvements to build housing. In a decentralized market economy with policy instruments that implement first-best outcomes, it must be that the price of the improvement input (numeraire) equals the marginal revenue product of improvements . As such, in any market allocation that can implement the first best, the real estate developer's FOC must be

$$1 = p_{Ri} \, \frac{\partial H_{Ri}}{\partial X_{Ri}} \, .$$

For a government with only land use policy instruments, any non-zero level of land use regulations $T_{Ri} \neq 1$ distorts the marginal revenue product of housing from improvements. The real-estate developer's first-order condition (FOC) in the presence of land use regulations is:

$$1 = T_{Ri} \, p_{Ri} \, \frac{\partial H_{Ri}}{\partial X_{Ri}}$$

For the two FOC to coincide, T_{Ri} must equal 1. $T_{Ri} = 1$ would however lead to a no-intervention market equilibrium which we know is inefficient.

The intuition behind why the land use regulations are not first-best is that land use regulations affect the quantity of *real estate* in each location, while urban externalities originate from *people*. By the "principle of targeting" (Bhagwati and Ramaswami, 1963; Johnson, 1965) , first-best policies should target the source of the externality (people), rather than a related good (housing). The location-based wage subsidies and taxes that Fajgelbaum and Gaubert (2020) derive analytically, which act directly on people, can implement first-best policy. In contrast, residential and commercial land use regulations $\{T_{Ri}, T_{Ci}\}$, which distort the real estate market, cannot perfectly internalize the urban externalities caused by the concentration of workers and residents.

4.3 Finding welfare-maximizing land use regulations using automatic differentiation

While residence- and workplace-specific taxes and subsidies in the spirit of Fajgelbaum and Gaubert (2020) can implement the first best, they are often not politically or administratively feasible. In practice, local governments commonly rely on second-best land-use regulations to address urban agglomeration and congestion externalities. This motivates computing the welfare-maximizing policy when the government's only instrument is land-use regulation.

There are no analytical expressions for the optimal land-use policy given the complex structure of the general equilibrium urban model. I therefore solve numerically for the local government's problem of choosing residential land-use regulations $\{T_{Ri}\}$ across locations i to maximize aggregate welfare, as outlined in (4.1).

I search for optimal land-use policies using gradient-based optimization, where the policy

gradients are obtained by implicit automatic differentiation of the general equilibrium system. All agents in the economy—workers, firms, and real-estate developers—take the chosen level of T_R as given, and their interactions produce an equilibrium outcome vector X that solves the system

$$F(X; T_R) = 0.$$

Accordingly, the endogenous equilibrium objects X can be written as functions of the policy vector, $X = X(T_R)$. The policymaker's problem can therefore be written as

$$\max_{T_R} \mathcal{O}(T_R), \qquad \mathcal{O}(T_R) \equiv O\big(X(T_R), T_R\big).$$

To solve for the optimal T_R, I use a two-step nested fixed-point algorithm in the spirit of Rust (1987), with the gradients used in the outer step obtained through automatic differentiation. In the inner step, for a given guess of the policy vector T_R, I solve for the equilibrium values $X(T_R)$ that satisfy

$$F\big(X(T_R),\, T_R\big) = 0. \tag{4.2}$$

The inner step is solved numerically using standard fixed-point and contraction-mapping methods. At the kth outer step, the candidate policy vector T_R^k is updated according to

$$T_R^{k+1} = T_R^k + \gamma\, H_k^{-1} \nabla_{T_R} \mathcal{O}(T_R^k),$$

where $\nabla_{T_R}\mathcal{O}$ is the gradient of the objective with respect to the policy vector T_R, and H_k^{-1} is the inverse Hessian of $\mathcal{O}$ that maps the gradient into a curvature-adjusted search direction.

A conventional way to compute the policy gradient would be to use numerical first differences. Under this approach, one perturbs each element of T_R in turn, resolves the equilibrium system $F(T_R) = 0$ for each perturbation to recover $X(T_R)$, and then recomputes the objective $\mathcal{O}$ to approximate the derivative numerically. Because T_R is high-dimensional, this

procedure requires solving the general equilibrium system many times at each outer-loop iteration, making it computationally infeasible in my setting. In my Cook County application, solving $F(T_R) = 0$ once takes approximately two minutes; T_R is a 175-dimensional policy vector, so even a single first-difference gradient update would require 175 equilibrium solves and around six hours of computation.

I therefore use automatic differentiation (AD) to obtain the relevant derivatives in a computationally tractable way. AD is a computationally efficient method for obtaining derivatives of functions and complex mappings, such as the large nonlinear general equilibrium system $F(\cdot)$, accurately and efficiently. AD decomposes a mapping such as $F(\cdot)$ into a sequence of elementary operations, such as addition, multiplication, and exponentiation, thereby constructing a computation graph. The chain rule is then applied systematically along this graph to propagate derivatives from inputs to outputs. In contrast to symbolic differentiation, AD does not manipulate algebraic expressions, and in contrast to finite-difference methods, it does not rely on numerical perturbations. Instead, it produces derivatives that are exact up to machine precision. Griewank and Walther (2008) and Margossian (2019) provide useful overviews of the methods and theory behind AD. Appendix A.5 provides a primer on automatic differentiation.

In my setting, AD must be combined with implicit differentiation. The reason is that the reduced-form objective $\mathcal{O}(T_R) = O(X(T_R), T_R)$ depends on T_R both directly and indirectly through the equilibrium mapping $X(T_R)$, which is not available in closed form and is instead defined implicitly by the equilibrium condition (4.2). Using the chain rule, the total derivative of $\mathcal{O}$ with respect to T_R is

$$\nabla_{T_R}\mathcal{O}(T_R) \;=\; \nabla_{T_R}O(X, T_R) \;+\; \nabla_X O(X, T_R)\,\frac{\mathrm{d}X}{\mathrm{d}T_R}.$$

Using the implicit function theorem to differentiate the inner equilibrium system $F(X(T_R), T_R) =$

0 gives

$$F_X(X, T_R)\,\frac{\mathrm{d}X}{\mathrm{d}T_R} \;+\; F_T(X, T_R) \;=\; 0 \qquad \Longrightarrow \qquad \frac{\mathrm{d}X}{\mathrm{d}T_R} \;=\; -\,F_X^{-1}(X, T_R)\,F_T(X, T_R).$$

Here, F_X and F_T are the Jacobians of F with respect to X and T_R, that can be obtained by automatic differentiation. Substituting this expression into the total derivative of $\mathcal{O}$ yields

$$\nabla_{T_R}\mathcal{O}(T_R) \;=\; O_T(X, T_R) \;-\; O_X(X, T_R)\,F_X^{-1}(X, T_R)\,F_T(X, T_R).$$

This expression gives the direction in which the policy vector T_R should be adjusted in the search for the welfare-maximizing policy.

The key computational bottlenecks in this expression are the Jacobian matrices F_X and F_T. If these derivatives were computed by first differencing, one would need to resolve the equilibrium system repeatedly to approximate each column of each Jacobian. In my setting, this would require on the order of $|X| + |T_R| = 5 \times 175 + 175$ equilibrium solves for each outer-loop update, which is computationally prohibitive.

Automatic differentiation avoids this bottleneck by exploiting the structure of the computation graph and reusing intermediate calculations. Once a node in the graph has been evaluated, its value can be stored and reused by downstream derivative calculations. This is the intuition behind the "cheap gradient principle" emphasized by Griewank and Walther (2008): the cost of computing the gradient of a scalar-valued function via AD is typically of the same order of magnitude as the cost of evaluating the function itself. In practice, this makes AD far more efficient than finite-difference methods in high-dimensional optimization problems such as mine.

Automatic differentiation is now foundational in machine learning, where it is routinely used to compute gradients for models with millions or billions of parameters. There are sev-

eral mature software packages for constructing computational graphs and evaluating derivatives, including TensorFlow, PyTorch, JAX, and CasADi. Sargent and Stachurski (2025) provides an excellent set of lectures for economists on implementing automatic differentiation methods in economic problems. Sargent and Yang (2025b) and Sargent and Yang (2025a) illustrate how automatic differentiation can be used to solve canonical Calvo and Ramsey policy problems.

Despite its widespread use in machine learning and optimal control, AD remains relatively uncommon in economics, especially in optimal policy and constrained optimization problems. A particularly relevant application is Wang et al. (2024), which uses AD to solve for optimal unilateral and multilateral tariff policies in a general equilibrium trade model. Their setting is closely analogous to mine: I solve for optimal land-use policy subject to equilibrium conditions in a quantitative spatial model, while they solve for optimal tariffs subject to equilibrium conditions in a quantitative trade model.

Another attractive feature of the implicit automatic-differentiation approach is its modularity. I code the general equilibrium system once and can reuse it for alternative planner objectives. Because the inner equilibrium condition $F(X; T_R) = 0$ is separate from the planner's objective $O(X, T_R)$, I can compute and cache AD graphs for F_X and F_{T_R} independently of those for O_X and O_{T_R}. I implement the solution by writing $F(X; T_R)$ and $O(X, T_R)$ as symbolic graphs from which sparse Jacobians and Hessians are generated and exploited automatically.

4.3.1 Implementation

To obtain the welfare–maximizing residential policy $T_R = \{T_{Ri}\}_i$, I use a nested fixed-point solution algorithm with automatic differentiation to compute gradients of the objective with-

out costly first-difference approximations. This is implemented in Python using CasADi[2]. I solve the optimization problem in logs because the general equilibrium system $F(X)$ is mostly multiplicative. Taking logs converts products, ratios, and powers into sums, differences, and scalar multiples, yielding a more compact and better-scaled system that is numerically more stable and faster to differentiate.

I map the gradient to a search direction using the quasi-Newton Broyden-Fletcher-Goldfarb-Shanno (BFGS) method, and enforce the bounds $0 \leq T_{Ri} \leq 1$ with the IPOPT interior-point optimizer when policy guesses approach the boundary. Automatic differentiation provides exact first derivatives of the objective, which are sufficient for first-order gradient-based search. Second-order Newton-type methods can converge faster by also using the Hessian matrix of second derivatives, but computing the full Hessian of second derivatives by automatic differentiation adds substantial computational cost in my high-dimensional setting. Instead, BFGS uses changes in gradients across iterations to construct an approximation to the Hessian, motivated by the multivariate mean value theorem. This yields a curvature-adjusted search direction that is substantially cheaper than forming the exact Hessian while retaining much of the speed advantage of Newton methods. Details on BFGS and IPOPT are in Appendices A.6.1 and A.6.2.

The combination of AD for first derivatives, BFGS for economical Hessian approximations, and IPOPT for feasibility near bounds yields stable performance at a modest computational cost; for 175 locations (and $5 \times 175 = 875$ equilibrium variables), the optimal land use regulation problem is solved in roughly 30 minutes. To check that the solution is not driven by initialization, I solve the problem from a range of diffuse, uniform starting guesses, setting $T_{Ri} \in \{0.4, 0.6, 0.8, 0.9\}$ for all i. Across all scenarios, the starting guess does not affect the welfare-maximizing policy that I find, which suggests that the computed solution

2. While AD is commonly implemented in ML packages such as PyTorch, TensorFlow, and JAX, CasADi is a symbolic–algorithmic differentiation and nonlinear-optimization framework that generates sparse Jacobians/Hessians and interfaces with IPOPT for bound constraints on T_R.

is a global maximum rather than a local optimum.

4.4 Optimal land use policies in Cook County

I compute optimal land use regulations $\{T_{Ri}\}$ across Cook County neighborhoods using my AD procedure and model fundamentals recovered by model inversion. I compare economic outcomes under optimal regulations against current regulations and a no regulation counterfactual with $T_{Ri} = 1$ across all locations. I first consider a no-externalities world, setting $\phi_{amen} = \phi_{prod} = 0$. I then examine optimal land use regulations under amenity and productivity externalities, varying the strength and sign of amenity spillovers to study their implications for the optimal policy.

4.4.1 No-externality case

I first examine the optimal shape of land use regulations without any urban congestion or agglomeration externalities; $\phi_{amen} = 0, \phi_{prod} = 0$. A natural prior is that, absent externalities, the planner would choose no regulation everywhere, i.e., $T_{Ri} = 1$ across all locations, as there are no obvious externalities to correct for.

This is not the case for optimal policy in my model. The optimal land use policy across locations features generally low levels of regulation, at an average $T_{Ri} = 0.87$. Panel A of Figure 4.1 plots the optimal distribution of residential land use regulations in the case without productivity and amenity spillovers. There are three clusters of regulatory stringency: in downtown, at the productive urban sub-center in western Cook County, and in the highest-income locations in the northern suburbs. Table 4.1 reports welfare, income, and land rent outcomes under the optimal policy, as compared to the baseline with the existing T_R^0 that I have recovered in Section 2.6.3.

The scope for government intervention, despite there being no explicit spillovers ($\phi_{amen} = \phi_{prod} = 0$), arises because idiosyncratic Fréchet preferences over location–workplace pairs $\epsilon_{ij}(\omega)$ generate a "virtual spillover" (Fajgelbaum and Gaubert, 2020, 2025). There exists a rationale for transfers under a quantitative spatial model with idiosyncratic utility shocks even without explicit spillovers. Under maximization of utilitarian ex-ante expected utility across all agents, average utilities in a residence location decline as marginal, lower-draw agents select in. As such, there is a redistributive motive to shift consumption toward low-amenity, low workplace market access locations whose residents have a higher idiosyncratic draw and, thus, have higher marginal utility of consumption. Because in my setting land income is pooled and rebated equally, tightening regulation (lower T_{Ri}) in locations with high housing prices, where amenities and workplace market access are high, raises the land revenue share. Land use regulations function like a place-based tax where the proceeds accrue to everyone, moving toward equalizing ex-ante expected utilities across space.

I compute optimal land use regulations in an alternative setting individuals receive no land rents at all; $\bar{y} = 0$ for all agents. In this no land rent distribution case, all land is owned by an absentee landlord who does not factor into the social planner's objective function. In this no redistribution case, all individual expenditures come only from wages. The computed optimal policy is $T^*_{Ri} = 1$ at every location, indicating that a no-regulation optimal land use policy when rents are not distributed.

Without the distribution of rents, regulation has no fiscal return; it only restricts real estate supply, bids up housing prices, and lowers consumption. Regulation is wasteful as the land rents created by regulation accrues to the absentee landlord. Any potential externality gains created by regulation are dominated by the welfare cost of reduced housing and consumption. To further examine the dynamics of the absentee landlord case, I solve for optimal regulations with the absentee landlord in the presence of urban exteranlities (ϕ_{amen}, ϕ_{prod} values used in the next section) . The no–regulation optimum holds in the presence of ur-

ban externalities when land rents are not distributed to individuals. To probe the absentee landlord case further, I solve for optimal regulations with $\bar{y} = 0$ while allowing for urban externalities, using the same $(\phi_{amen}, \phi_{prod})$ values as section 4.4.2. The computed optimum remains no regulation: $T^*_{Ri} = 1$ for all i when land rents are not redistributed.

I rebate land rents uniformly to residents to keep rents inside the economy. I treat agents as ex-ante identical and do not track identities; a common lump-sum rebate respects this anonymity and avoids introducing wealth heterogeneity and ownership considerations that, while interesting, are separate from my spatial problem of considering the ability of land use regulations to address urban externalities. [3]

4.4.2 Optimal regulations under urban externalities

I now examine the optimal land use policy in the presence of urban externalities. I adopt a productivity spillover elasticity of $\phi_{prod} = 0.07$ based on within-city estimates in Ciccone and Hall (1996) and Ahlfeldt et al. (2015). Positive productivity externalities reflect what is commonly called the "fundamental tradeoff of urban economics"– cities exist because of urban agglomeration forces, despite the costs of concentrating people within the city.

The relationship between amenities and residential population is ambiguous because opposing mechanisms operate simultaneously. The concentration of residents may increase amenities by supporting a greater variety of services and enabling more efficient provision of public goods. At the same time, population density can erode amenities through congestion, competition for public space, increased crime, and local pollution. I examine optimal regulations under three values of amenity spillover elasticity: $\phi_{amen} \in \{0.15, -0.04, -0.15\}$. The positive value follows Ahlfeldt et al. (2015), while the negative values capture amenity externalities arising from congestion (Duranton and Puga, 2023) and crime (Glaeser and

3. An alternative setting that can be explored is the local ownership of land by residents.

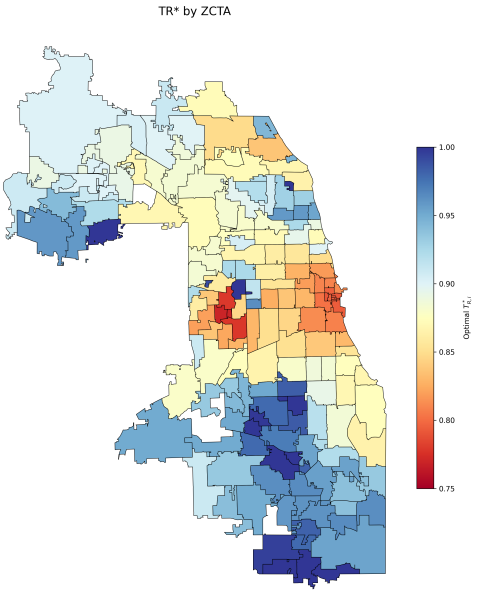

(a) $\phi_{amen} = 0$, $\phi_{prod} = 0$; no spillovers case

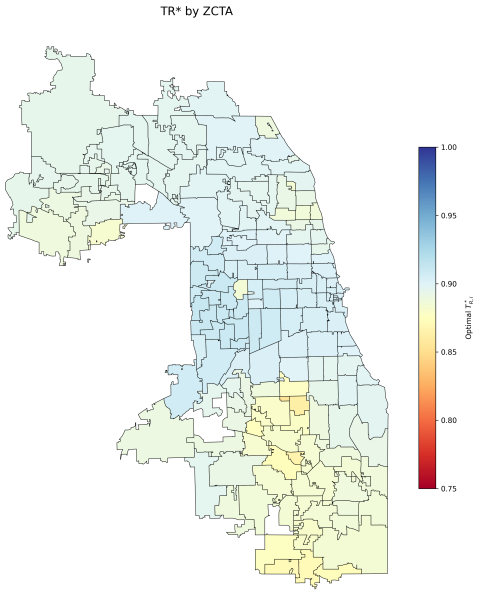

(b) $\phi_{amen} = 0.15$, $\phi_{prod} = 0.07$; positive amenity spillovers. High contrast version in Figure 4.2

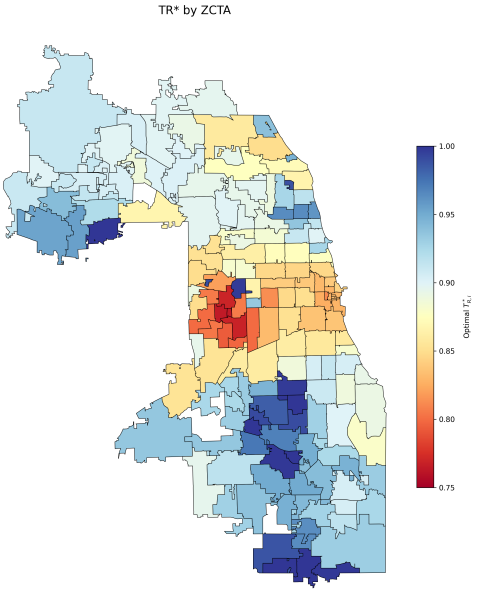

(c) $\phi_{amen} = -0.04$, $\phi_{prod} = 0.07$; weak negative amenity spillovers

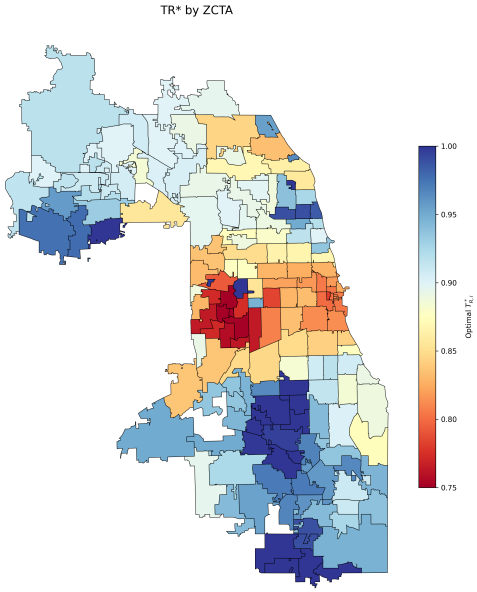

(d) $\phi_{amen} = -0.15$, $\phi_{prod} = 0.07$; strong negative amenity spillovers

All values are plotted on the same color scale.

Figure 4.1: Optimal T_R^* across different amenity and productivity spillovers ϕ_{amen}, ϕ_{prod}.

Sacerdote, 1999).

I compare regulations, aggregate welfare, wages, and land rents under current observed values of regulation (T_R^0), a no regulation counterfactual (all $T_{Ri} = 1$), and optimal land use policy T_R^*. Table 4.1 reports economic outcomes under the different values of ϕ_{amen}. Figure 4.1 plots the distribution of optimal land use regulations T_R^* across space under the different scenarios.

Across all spillover assumptions, optimal regulation T_{Ri}^* delivers welfare gains of around 2% relative to existing baseline regulations. By contrast, setting all $T_{Ri} = 1$ yields gains that are highly sensitive to amenity externalities: 1.5% when amenity spillovers are positive ($\phi_{amen} = 0.15$), 0.24% with mildly negative amenities ($\phi_{amen} = -0.04$), and only 0.05% when amenities are strongly negative ($\phi_{amen} = -0.15$). This decline with more negative amenity spillovers is because full deregulation over-concentrates residents in the previously regulated northern and western suburbs. Congestion in amenities, which is stronger when ϕ_{amen} is more negative, attenuates the welfare gains from increased housing productivity under deregulation.

Across all spillover assumptions considered, the welfare-maximizing regulation T_{Ri}^* yields gains of roughly 2% relative to existing baseline regulations. In contrast, setting $T_{Ri} = 1$ (full deregulation) produces gains that vary sharply with amenity spillovers: 1.5% when $\phi_{amen} = 0.15$, 0.24% when $\phi_{amen} = -0.04$, and just 0.05% when $\phi_{amen} = -0.15$. The declining payoff of deregulation as ϕ_{amen} becomes more negative arises because full deregulation concentrates residents in the formerly regulated northern and western suburbs. Under negative amenity spillovers, this concentration worsens congestion and reduces local amenities, attenuating the welfare benefits from full deregulation due to higher housing productivity. Optimal regulations yield strictly higher welfare than full deregulation[4], with especially

4. Full deregulation($T_{Ri} = 1$ everywhere) is within the feasible set of optimal regulations, so welfare under computed optimal regulations must be weakly larger than welfare under full deregulation.

	ΔWelfare (rel. to baseline)	Agg./ Min. T_{Ri}	Ave. Wage ($)	Land rents pp. $\bar{y}$ ($)
$\phi_{amen} = 0$, $\phi_{prod} = 0$; no spillovers				
Baseline T_R^0		0.858/0.541	72506	9602
All $T_{Ri} = 1$	+0.23%	1/1	74114	6403
Optimal T_R^*	+1.4%	0.871/0.769	72306	8839
$\phi_{amen} = 0.15$, $\phi_{prod} = 0.07$; Positive amenity and productivity externalities				
Baseline T_R^0		0.858/0.541	72501	10808
All $T_{Ri} = 1$	+1.5%	1/1	76529	7475
Optimal T_R^*	+2.1%	0.897/0.853	74132	9836
$\phi_{amen} = -0.04$, $\phi_{prod} = 0.07$; Negative amenity and productivity externalities				
Baseline T_R^0		0.858/0.541	73587	11002
All $T_{Ri} = 1$	+0.24%	1/1	76717	7491
Optimal T_R^*	+1.5%	0.875/0.764	73376	10099
$\phi_{amen} = -0.15$, $\phi_{prod} = 0.07$; Strong negative amenity and productivity externalities				
Baseline T_R^0		0.858/0.541	73797	11029
All $T_{Ri} = 1$	+0.05%	1/1	76815	7500
Optimal T_R^*	+1.6%	0.862/0.717	73472	10277

The baseline T_R^0 case refers to the economic outcomes under current existing regulations for Cook County. The all $T_{Ri} = 1$ case refers to the counterfactual where there is minimal regulation across all locations; $T_{Ri} = 1 \ \forall i$. Welfare in the optimum is calculated as a percentage change relative to the existing baseline regulatory scheme T_R^0. Aggregate T_{Ri} refers to city-level wedges, aggregated from ZIP-level T_{Ri}.

Table 4.1: Welfare, average regulation T_{Ri} across locations, average wages, and per-person land-rents across different scenarios

significant gains under negative amenity externalities. The presence of amenity and productivity externalities provides a clear rationale for non-zero welfare-maximizing land-use regulations.

I consider the notion of aggregate price wedge $T_{R,\text{city}}$ across the city, using the same logic that motivates the aggregation of parcel-level τ_ω to neighborhood level T_{Ri}. The city-level regulatory wedge is

$$T_{R,\text{city}} = \left(\frac{\sum_i X_i / T_{Ri}}{\sum_i X_i / T_{Ri}^{\frac{1}{1-\gamma}}} \right)^{\frac{1-\gamma}{\gamma}}.$$

Aggregate city-level regulatory stringency across all values of ϕ_{amen} is lower compared to the baseline T_R^0, as shown in column 2 of Table 4.1. In every scenario, there is also overall deregulation in optimum. However, due to the attenuation from congestion , optimal aggregate regulatory stringency is lowest with strong negative amenity spillovers and highest with strong positive spillovers.

The spatial pattern of optimal regulation is similar across the three cases where amenity spillovers are either negative or zero ($\phi_{amen} = 0\,, -0.04\,, -0.15$). The distribution of T_{Ri}^* across these three scenarios is plotted in panels (A), (C), and (D) of Figure 4.1. The no-externality scenario features optimal regulations similar to the negative-amenity-spillover case due to the Fajgelbaum and Gaubert "virtual spillover" logic as explained in Section 4.4.1.

In the negative amenity spillover case, optimal land use regulations are more stringent (lower T_{Ri}^*) around downtown, the western suburbs, and the southern suburbs. There are relatively more permissive regulations (higher T_{Ri}^*) on the South Side. Downtown, the western suburbs, and the northern suburbs have good access to high-paying jobs. As I discuss further in Section 4.4.4, the planner addresses overcrowding stemming from negative amenity spillovers in these attractive locations by setting stringent land use regulations. Shifting

ϕ_{amen} changes the average stringency of regulations but not the pattern of optimal regulations across space. The relative distribution of regulation across space in optimum looks somewhat similar to the relative distribution of *current* regulations, as plotted in Figure 2.2. A notable departure is downtown Chicago, where current rules are very permissive but the optimal policy is stringent.

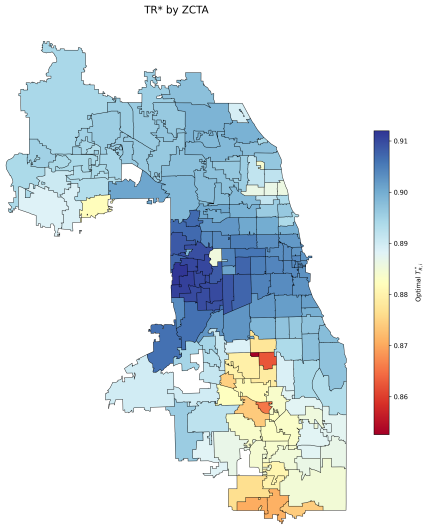

Figure 4.2: Optimal land use regulations T_{Ri} across space, $\phi_{amen} = 0.15, \phi_{prod} = 0.07$

The pattern of optimal land use regulations across space is different when amenity spillovers are positive ($\phi_{amen} = 0.15$). This is plotted in Figure 4.2. The aggregate city-level $T_{R,\text{city}}$ is highest among all cases at 0.897. Given positive externality spillovers, the policymaker would want to concentrate, rather than spread out, people in the high productivity and high resident market access areas of downtown and the northern and western suburbs. This is done by setting relatively high levels of T_{Ri}, and thus low levels of regulatory stringency in these areas.

Figure 4.3 plots changes in worker and resident populations when replacing current regulations T_R^0 with the optimal policy T_R^* under both positive ($\phi_{amen} = 0.15$) and negative ($\phi_{amen} = -0.15$) amenity spillovers. For both signs of amenity spillovers, north and northwest Cook County gain workers and residents at the expense of the south. Worker gains are

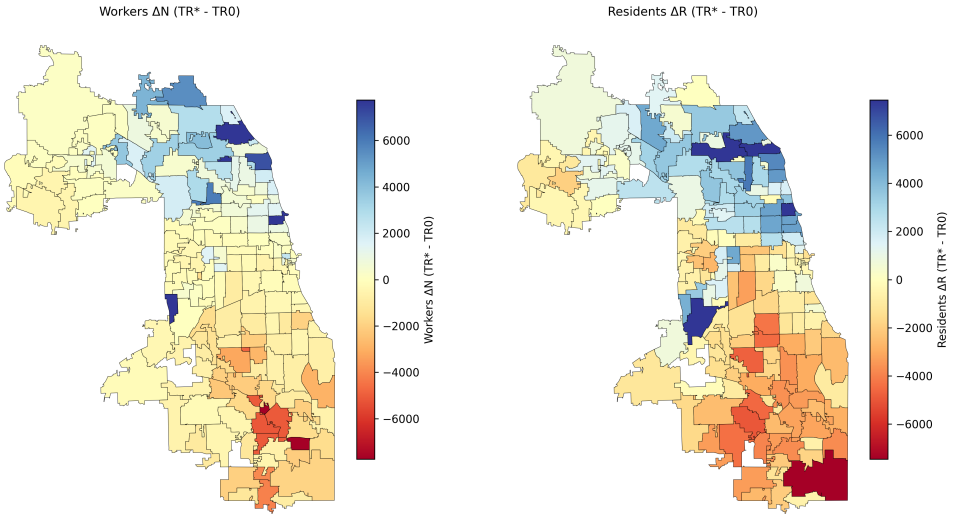

(a) Workers(left) and residents(right) changes, positive amenity spillovers $\phi_{amen} = 0.15$

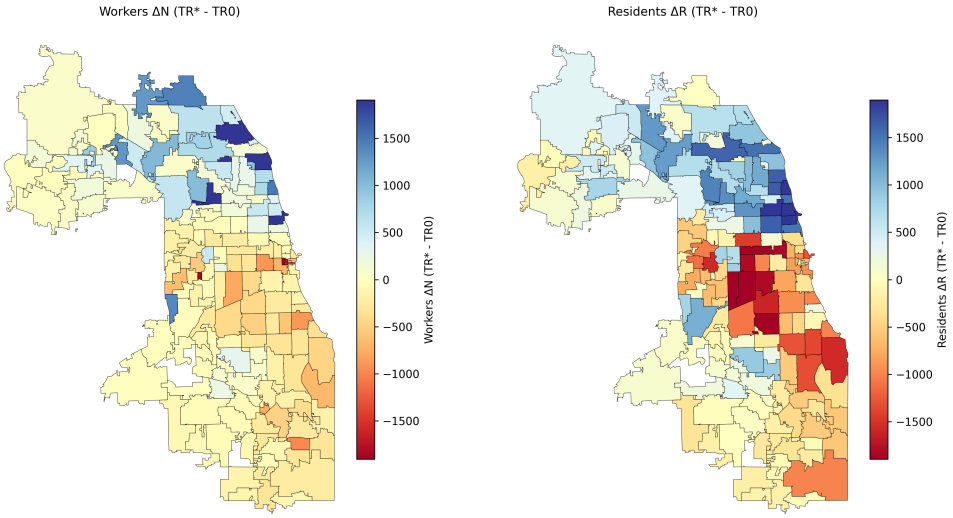

(b) Workers(left) and residents(right) changes, negative amenity spillovers: $\phi_{amen} = -0.15$

Figure 4.3: Worker and resident population changes, optimal regulations T_R^* vs. current regulations T_R^0

most pronounced in the most productive northern neighborhoods in both cases. The qualitative pattern of south to north movement of workers and residents under optimal regulation does not depend on the sign of amenity spillovers.

The invariance of population changes to the sign of amenity spillovers is because al-

though under negative amenity spillovers the north remains *relatively* more regulated than other areas, the optimum still relaxes regulations in absolute *levels*. Under positive amenity spillovers, the north is deregulated both relatively and in levels. The interplay between relative and absolute levels of regulation also explains the smaller increase in workers and residents in the north under negative amenity spillovers (Panel B, Figure 4.3), as compared to the large increase in the positive amenity spillover case (Panel A).[5]

To isolate how the optimal policy across *space* affects the distribution of people, I compare the distribution of residents and workers under T_R^* with a constructed counterfactual featuring uniform regulations in every location at the level equal to the *aggregate* level of T_R^*. This thus compares T_R^* against a world where aggregate regulations are the same but there is no spatial variation in regulations across locations, isolating the spatial effect of optimal policy.

Figure 4.4 plots the changes in workers and residents under T_R^* compared to the counterfactual world with uniform regulation of the same aggregate stringency. Residential land use regulations mainly change the distribution of residents rather than workers. This is because changes in residential land use regulations affect worker populations indirectly through commuting costs and the wage adjustments needed to attract labor .

Panel A of Figure 4.4 plots the case when amenities rise with local population ($\phi_{amen} = 0.15$). The optimal policy concentrates people in the high downtown, northern suburbs, and western suburbs. This is because residents in these locations have high access to jobs (RMA as explained in Section 3.5) and the planner will want to concentrate people to take advantage of the positive amenity spillovers. The planner moves residents from the South Side, which has low access to jobs, to the high RMA downtown, west, and north.

Panel B of Figure 4.4 plots the case when additional residents reduce local amenities

5. The maximum increase in residents and workers is around 1500 for the negative amenity spillover case and 6000 for the positive case.

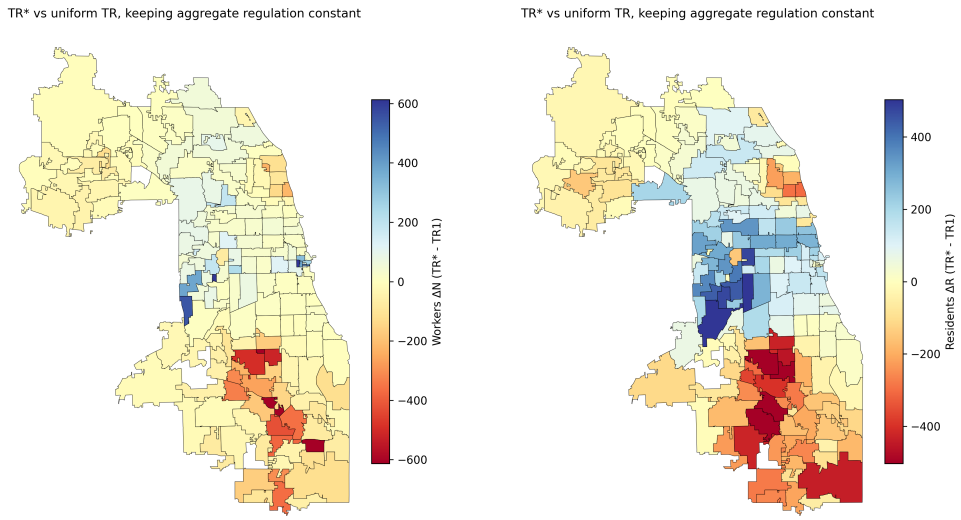

(a) Workers(left) and residents(right) changes, positive amenity spillovers $\phi_{amen} = 0.15$

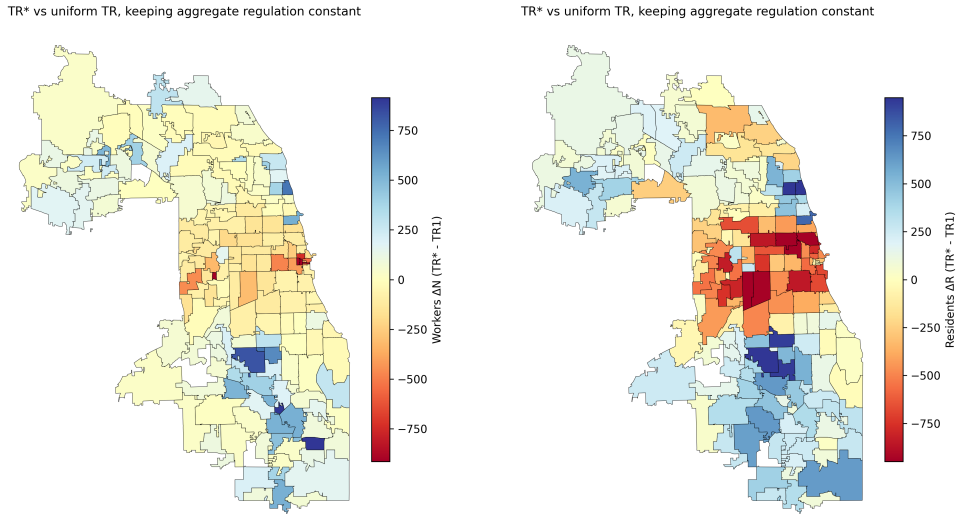

(b) Workers(left) and residents(right) changes, negative amenity spillovers: $\phi_{amen} = -0.15$

Figure 4.4: Worker and resident population changes, T_R^* vs. uniform T_R keeping aggregate city-level regulations constant

($\phi_{amen} = -0.15$). Productivity spillovers still reward some concentration in productive downtown areas, but the planner now spreads residents out, shifting mass away from high-wage downtown and northern locations and toward the lower-productivity south. This is because locations with high access to jobs will be overcrowded with residents, requiring more stringent land use regulations to address the negative amenity spillovers from residential

density.

Overall, baseline productivity differences anchor where the population is directed under T_R^*, while the sign of ϕ_{amen} governs how costly it is to stack residents in high RMA neighborhoods. The resulting north–south gradient in population changes reflects the spatial profile of optimal regulation.

4.4.3 The effect of regulations on local housing markets

The left panel of Figure 4.5 shows the change in the neighborhood-level regulatory price wedge under the optimal policy relative to its baseline level, in the $\phi_{amen} = -0.15$ calibration. The right panel shows the corresponding percent change in total housing floorspace. Neighborhoods that experience a relaxation of land use regulations under the optimal policy exhibit an expansion in residential floorspace. This pattern follows directly from the housing supply equation; land use regulations enter as a supply shifter (equivalently as a productivity shifter) for housing developers, so a reduction in the regulatory wedge raises effective housing productivity and increases equilibrium construction.

Deregulated neighborhoods also experience declines in per-square-foot housing prices, reflecting the outward shift in housing supply (left panel of Figure 4.6). In contrast, neighborhoods where the optimal policy tightens regulations display the opposite pattern: housing quantities contract and prices rise. These results formalize a common political concern surrounding deregulation—namely, that expanding housing supply may depress local property prices and thereby generate homeowner opposition.

The effect on total housing values—defined as the product of price and quantity—is theoretically ambiguous. Deregulation increases housing quantities but lowers prices, so the net effect depends on the relative magnitudes of these forces. The right panel of Figure 4.6 shows that, in general equilibrium, total housing values rise in deregulated neighborhoods and

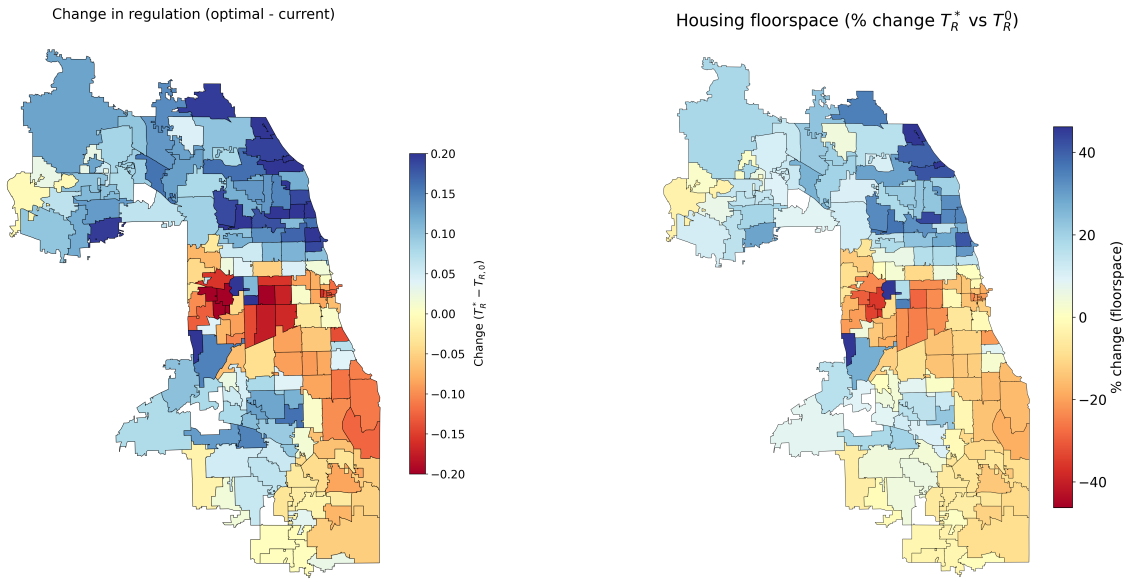

Figure 4.5: Changes in neighborhood regulatory wedge and housing quantity (%) under optimal regulations T_R^*, relative to baseline T_R^0.

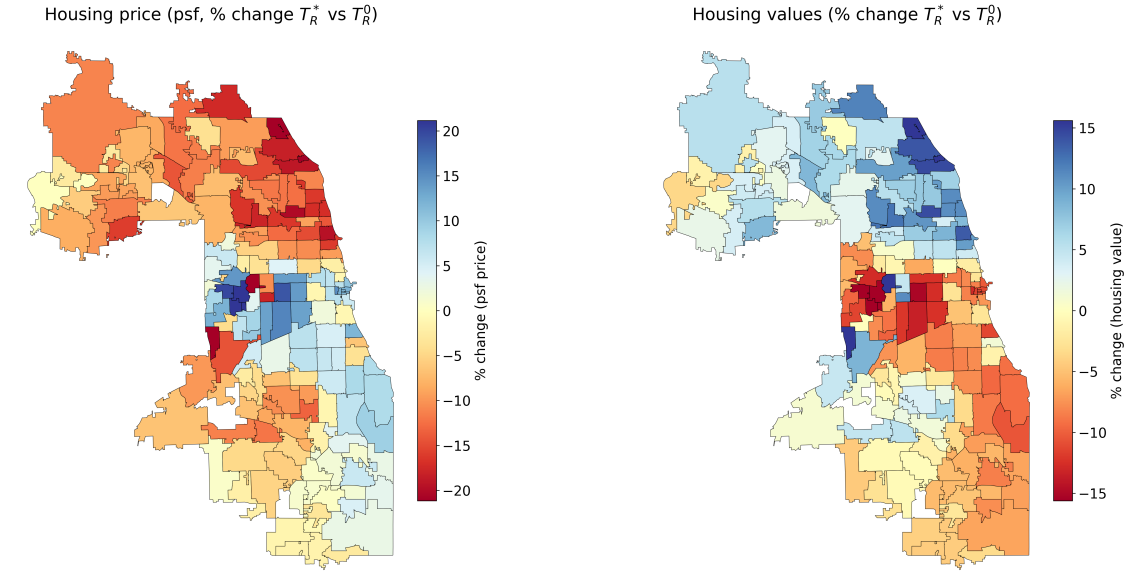

Figure 4.6: Spatial changes (%) in housing prices and housing values under optimal regulations T_R^*, relative to baseline T_R^0.

fall where regulations are tightened. This outcome contrasts with the common presumption that deregulation necessarily reduces local housing wealth.

The key mechanism is endogenous household relocation. Deregulated neighborhoods become more attractive because they offer a larger housing stock at lower prices, drawing in additional residents. The resulting increase in local demand mitigates the initial downward pressure on prices. In equilibrium, the expansion in housing quantities more than offsets the decline in per-square-foot prices, leading to higher total housing values in deregulated areas.

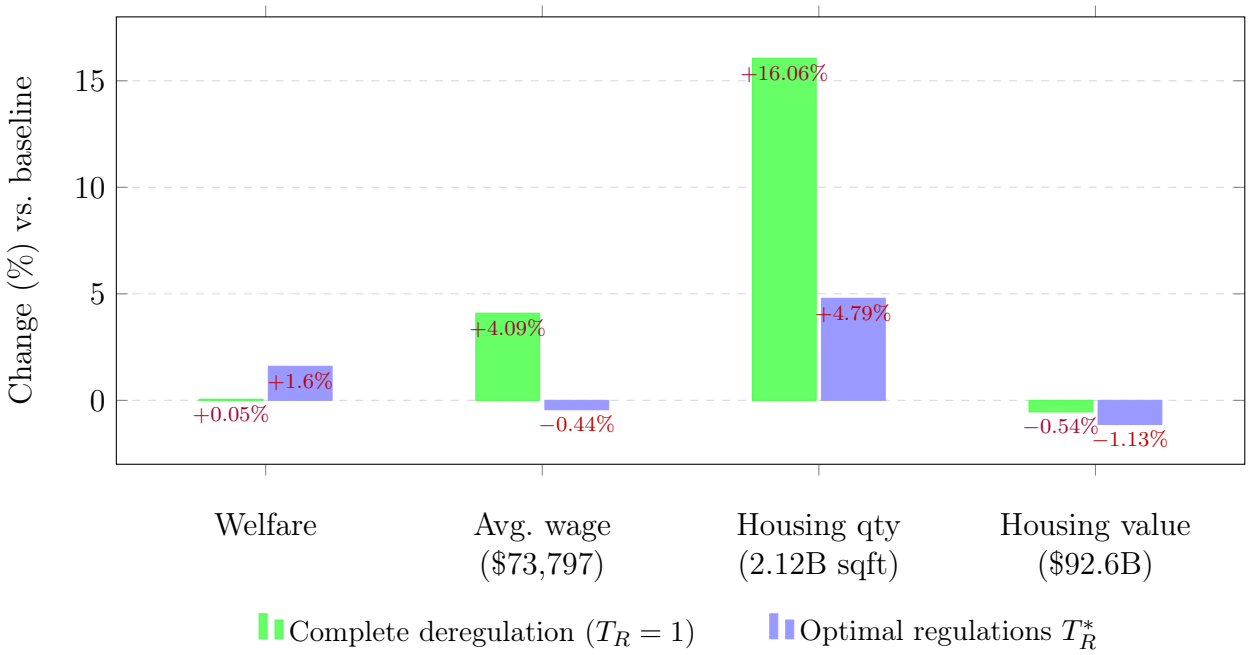

Figure 4.7: Aggregate outcomes under complete deregulation and optimal regulation, % change relative to baseline T_R^0.

At the county level, total housing stock in Cook County increases by 4.79% under the welfare-maximizing regulatory scheme, relative to a baseline of 2.12 billion square feet. At the same time, aggregate housing value declines by 1.13%. Under complete deregulation ($T_{Ri} = 1$ for all i), total housing stock would rise by a substantially larger 16.06%, while aggregate housing value would fall by 0.54%.

Thus, although complete deregulation generates a much larger expansion in construction, it delivers smaller aggregate welfare gains (Figure 4.7). In contrast, the optimal policy achieves more modest growth in total floorspace while better balancing aggregate welfare

improvements against distributional and political considerations.

4.4.4 What determines optimal land use regulations?

To examine the determinants of the level of optimal land use regulation, I examine the relationship between ZIP code level land use regulation T^*_{Ri} and local fundamental amenities, productivities, residential market access, and firm market access. Residential and firm market access are terms that summarize the spatial linkages across locations (Redding, 2024b). Residential market access (RMA) at location i is defined by

$$\text{RMA}_i = \sum_j \left(\frac{e_j}{d_{ij}} \right)^{\epsilon}.$$

RMA is a measure of the accessibility of high paying jobs e_j to residents living at location i, weighted by the commuting costs d_{ij} from home i to workplace j. Keeping amenities and housing prices constant, a residential location with high RMA would see more residents due to better accessibility to jobs in productive locations. RMA was previously used in equation (3.24) to recover local amenities.

Firm market access (FMA) calculates the access of firms to workers. A firm at workplace location j with high firm market access is close to residential locations i with a high number of residents (due to high amenities B_i and low real-estate prices p_{Ri}). FMA is given by

$$\text{FMA}_j = \sum_i \left(\frac{B_i}{d_{ij} p_{Ri}^{1-\alpha}} \right)^{\epsilon}.$$

Keeping wages constant, a workplace location with high FMA would see a higher labor supply due to its proximity to high-amenity, low rent locations.

Table 4.2 reports the results of a regression of the log of optimal locational land use

Dependent variable: $\log(T^*_{Ri})$				
	(1)	(2)	(3)	(4)
Spillovers:	$\phi_{amen} = 0.15$	$\phi_{amen} = 0.15$	$\phi_{amen} = -0.15$	$\phi_{amen} = -0.15$
log(productivity)	-0.0032**	-0.0031**	0.0096	0.0003
	(0.0016)	(0.0014)	(0.0118)	(0.0099)
log(amenity)	-0.0077***	0.0048***	0.0229***	-0.0266***
	(0.0015)	(0.0018)	(0.0076)	(0.0074)
log(resident MA)		0.0707***		-0.4590***
		(0.0077)		(0.0499)
log(firm MA)		-0.0658***		0.3756***
		(0.0105)		(0.0715)
R^2	0.142	0.448	0.050	0.436
N	175	175	175	175

Dependent variable: $\log(T^*_{Ri})$. Robust standard errors in parentheses. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4.2: Regression of log of ZIP-code optimal regulation T^*_{Ri} on amenities, productivities, resident MA, and firm MA

regulations T^*_{Ri} , against a combination of productivity, amenity, RMA and FMA, for both the case of both positive ($\phi_{amen} = 0.15$) and negative ($\phi_{amen} = -0.15$) amenity externalities.

Amenities and productivity alone explain little of the optimal policy, as evidenced by the low R^2 of columns (1) and (3) of Table 4.2. The addition of resident and firm market access terms (columns (2) and (4)) substantially increases the explanatory power of the regression ($R^2 = 0.448$ and 0.436). The elasticities of optimal T_{Ri} with respect to RMA and FMA in columns (2) and (4) are an order of magnitude higher than the elasticities for amenities and productivity. Given the *separation* of residences and workplaces due to commuting, *spatial linkages* (as summarized by RMA and FMA) determine land use optimal regulations more than *local* amenities and productivity.

When amenity spillovers are positive ($\phi_{amen} = 0.15$), higher residential market access and higher amenities are associated with less stringent land use regulations. Given positive amenity spillovers, the policymaker would want to concentrate residents in desirable locations with high RMA and high amenities by setting higher T_{Ri}.

When amenity spillovers are negative ($\phi_{amen} = -0.15$), the movement of optimal regulations with amenities, RMA, and FMA reverses. When the concentration of people reduces local amenities, the policymaker sets out to prevent overcrowding in desirable locations (high amenity and RMA) by implementing more stringent land use regulations in the form of lower T_{Ri}.

The movement of optimal residential regulation with *firm* market access is more subtle. Locations with high FMA already enjoy abundant labor supply, so the marginal productivity gain from adding another resident to increase labor supply is limited. With positive amenity spillovers, placing additional residents at these high-FMA, high–labor-supply job centers carries a high opportunity cost, as it forgoes large amenity externalities the same residents would generate in neighborhoods with high amenities and RMA. Therefore, the planner tightens regulation (lowering T^*_{Ri}) at high-FMA locations. When amenity spillovers are negative, the logic reverses: housing workers closer to firms helps relieve crowding in high-amenity, high-RMA areas, so the planner relaxes regulation at high-FMA job centers.

In general equilibrium, spatial linkages between locations are more important than local fundamentals in determining optimal regulations. The spatial pattern of optimal regulation depends *sign* of amenity externalities, underscoring the need to accurately estimate amenity spillovers.

CHAPTER 5

CONCLUSION

This dissertation studies optimal land use regulations within the city in the presence of amenity and productivity urban externalities. To model land use regulations in a micro-founded and model-consistent manner, I use a price wedge measure of regulatory stringency that is both recoverable at the parcel level and can be aggregated to any spatial unit. The price wedge measure of regulation can be recovered via property prices and building costs, which are available from property records and construction industry data. I estimate the regulatory wedge for Chicago and Cook County. The regulatory wedge aligns with general perceptions of regulations in Cook County and reflects features of Chicago's statutory land use zoning scheme.

I show theoretically that land use regulations cannot perfectly internalize urban externalities, which are generated by people rather than real estate. Land use regulations are therefore second-best and I solve for optimal regulations numerically. To handle the high-dimensional numerical problem, I develop a solution strategy that leverages automatic differentiation in a nested fixed-point framework. I compute optimal land use regulation across neighborhoods in Cook County for a range of strengths and signs of amenity spillovers.

The optimal land-use policy yields welfare gains of around 2% over current regulations and significantly exceeds the gains from proposed complete deregulation policies. Urban agglomeration externalities are fundamental to city formation (Fujita and Thisse, 2002; Glaeser and Gottlieb, 2009); absent such forces, there would be little reason for firms and households to form urban clusters. While current regulations are often too stringent, eliminating land use regulations entirely ignores urban externalities, making the case for *optimal* land use policy. In Cook County, existing regulations in the affluent northern suburbs are excessively stringent; across all parameterizations of urban externalities, the optimal policy reallocates

residents and employment from the south to the north of Cook County.

The spatial distribution of optimal regulations hinges on the sign of amenity spillovers. If residential density raises amenities, the policymaker concentrates residents in desirable neighborhoods with high residential market access and high amenities; if density lowers amenities, the policymaker disperses residents in the same locations to avoid overcrowding.

The dependence of the optimal policy on the sign and magnitude of urban externalities highlights the need for accurate estimation. While I assume constant spillover elasticities for amenities and productivity across neighborhoods, spillovers are plausibly location-specific: downtown areas often exhibit productivity and amenity gains from density, whereas increased density may cause congestion, pressure on public services, or perceived neighborhood change in some suburbs. Future work should estimate location-specific spillovers and allow for heterogeneous externalities, so that land use regulation encourages housing supply where gains from density are large and restrains it where the "demons of density"[1] reign.

1. Glaeser (2011).

REFERENCES

Ahlfeldt, Gabriel and Daniel McMillen. 2014. "New Estimates of the Elasticity of Substitution of Land for Capital." ERSA conference papers, European Regional Science Association.

Ahlfeldt, Gabriel M. and Daniel P. McMillen. 2018. "Tall Buildings and Land Values: Height and Construction Cost Elasticities in Chicago, 1870–2010." 100 (5): 861–875. 10.1162/rest_a_00734.

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2015. "The Economics of Density: Evidence from the Berlin Wall." *Econometrica* 83 (6): 2127–2189. 10.3982/ECTA10876.

Ahluwalia, Gopal, Kermit Baker, and Kent W. Colton. 2022. "Concentration in the Homebuilding Industry: Trends, Strategies, and Prospects." Tech. rep., Joint Center for Housing Studies of Harvard University. https://www.jchs.harvard.edu/sites/default/files/research/files/harvard_jchs_homebuilding_industry_concentration_ahluwalia_2022.pdf.

Allen, Treb, Costas Arkolakis, and Xiangliang Li. 2016. "Optimal City Structure." Working paper, Yale University.

Babalievsky, Fil, Kyle F. Herkenhoff, Lee E. Ohanian, and Edward C. Prescott. 2023. "The Impact of Commercial Real Estate Regulations on U.S. Output." NBER Working Paper 31895, National Bureau of Economic Research. 10.3386/w31895.

Barr, Jason and Jeffrey P. Cohen. 2014. "The Floor Area Ratio Gradient: New York City, 1890–2009." *Regional Science and Urban Economics* 48: 110–119. 10.1016/j.regsciurbeco.2014.03.004.

Bastani, Houtan and Luca Guerrieri. 2008. "On the Application of Automatic Differentiation to the Likelihood Function for Dynamic General Equilibrium Models." International Finance Discussion Papers 920, Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/pubs/ifdp/2008/920/.

Baum-Snow, Nathaniel and Lu Han. 2024. "The Microgeography of Housing Supply." *Journal of Political Economy* 132 (6): 1897–1946. 10.1086/728110.

Berry, Christopher. 2025. "An Evaluation of Progress on Residential Assessment Fairness in Cook County." Tech. rep., University of Chicago, Mansueto Institute for Urban Innovation and Center for Municipal Finance, Chicago, IL. https://propertytaxproject.uchicago.edu/files/2025/09/Kaegi-Evaluation-9_04.pdf. This draft: September 3, 2025.

Bhagwati, Jagdish N. and V. K. Ramaswami. 1963. "Domestic Distortions, Tariffs and the Theory of Optimum Subsidy." *Journal of Political Economy* 71 (1): 44–50. 10.1086/258733.

Boehl, Gregor. 2023. "HANK on Speed: Robust Nonlinear Solutions using Automatic Differentiation." https://ssrn.com/abstract=4433585. SSRN Working Paper No. 4433585.

Brueckner, Jan K. 2007. "Urban Growth Boundaries: An Effective Second-Best Remedy for Unpriced Traffic Congestion?" *Journal of Housing Economics* 16 (3–4): 263–273.

Brueckner, Jan K., Shihe Fu, Yizhen Gu, and Junfu Zhang. 2017. "Measuring the Stringency of Land Use Regulation: The Case of China's Building Height Limits." 99 (4): 663–677. 10.1162/REST_a_00650.

Brueckner, Jan K. and Ruchi Singh. 2020. "Stringency of land-use regulation: Building heights in US cities." *Journal of Urban Economics* 116: 103239. 10.1016/j.jue.2020.103239.

Cai, Hongbin, Zhi Wang, and Qinghua Zhang. 2017. "To Build Above the Limit? Implementation of Land Use Regulations in Urban China." *Journal of Urban Economics* 98: 223–233. 10.1016/j.jue.2016.03.003.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan. 2007. "Business Cycle Accounting." *Econometrica* 75 (3): 781–836. 10.1111/j.1468-0262.2007.00768.x.

Ciccone, Antonio and Robert E. Hall. 1996. "Productivity and the Density of Economic Activity." *American Economic Review* 86 (1): 54–70.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2015. *The Economics of Agglomeration: Cities, Industrial Location, and Globalization*. Princeton, NJ: Princeton University Press, 2nd ed.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2021. "The Production Function for Housing: Evidence from France." *Journal of Political Economy* 129 (10): 2766–2816. 10.1086/715076.

D'Amico, Leonardo and Maxim Alekseev. 2025. "Capital Market Integration and Growth Across the United States." Tech. rep. https://leonardodamico.github.io/papers/damico_jmp.pdf. Working Paper.

Davis, Morris A, William D Larson, Stephen D Oliner, and Jessica Shui. 2021. "The price of residential land for counties, ZIP codes, and census tracts in the United States." *Journal of Monetary Economics* 118: 413–431. 10.1016/j.jmoneco.2020.12.005.

Davis, Morris A. and François Ortalo-Magné. 2011. "Household Expenditures, Wages, Rents." *Review of Economic Dynamics* 14 (2): 248–261. 10.1016/j.red.2010.11.003.

Desmet, Klaus and Esteban Rossi-Hansberg. 2013. "Urban Accounting and Welfare." *American Economic Review* 103 (6): 2296–2327. 10.1257/aer.103.6.2296.

Duranton, Gilles and Diego Puga. 2015. "Urban Land Use." In *Handbook of Regional and Urban Economics*, vol. 5A, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, pp. 467–560. Amsterdam: Elsevier.

Duranton, Gilles and Diego Puga. 2023. "Urban Growth and Its Aggregate Implications." *Econometrica* 91 (6): 2219–2259.

Eaton, Jonathan and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70 (5): 1741–1779.

Epple, Dennis, Brett Gordon, and Holger Sieg. 2010. "A New Approach to Estimating the Production Function for Housing." *American Economic Review* 100 (3): 905–925. 10.1257/aer.100.3.905.

Fajgelbaum, Pablo D. and Cécile Gaubert. 2020. "Optimal Spatial Policies, Geography, and Sorting." *The Quarterly Journal of Economics* 135 (2): 959–1036.

Fajgelbaum, Pablo D. and Cécile Gaubert. 2025. "Optimal Spatial Policies." NBER Working Paper 33493, National Bureau of Economic Research.

Fujita, Masahisa and Jacques-François Thisse. 2002. *Economics of Agglomeration: Cities, Industrial Location, and Regional Growth*. Cambridge: Cambridge University Press. 10.1017/CBO9780511805660.

Ganong, Peter and Daniel Shoag. 2017. "Why has regional income convergence in the U.S. declined?" *Journal of Urban Economics* 102: 76–90. 10.1016/j.jue.2017.07.002.

Glaeser, Edward L. 2011. *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier*. New York: Penguin Press.

Glaeser, Edward L. and Joshua D. Gottlieb. 2009. "The Wealth of Cities: Agglomeration

Economies and Spatial Equilibrium in the United States." *Journal of Economic Literature* 47 (4): 983–1028. 10.1257/jel.47.4.983.

Glaeser, Edward L. and Joseph Gyourko. 2018. "The Economic Implications of Housing Supply." *Journal of Economic Perspectives* 32 (1): 3–30.

Glaeser, Edward L and Bruce Sacerdote. 1999. "Why is there more crime in cities?" *Journal of Political Economy* 107 (S6): S225–S258.

Griewank, Andreas and Andrea Walther. 2008. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. Philadelphia: SIAM, 2 ed. 10.1137/1.9780898717761.

Gyourko, Joseph, Jonathan S. Hartley, and Jacob Krimmel. 2021. "The local residential land use regulatory environment across U.S. housing markets: Evidence from a new Wharton index." *Journal of Urban Economics* 124. 10.1016/j.jue.2021.103337.

Gyourko, Joseph and Raven Molloy. 2015. "Chapter 19 - Regulation and Housing Supply." In *Handbook of Regional and Urban Economics*, vol. 5 of *Handbook of Regional and Urban Economics*, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, pp. 1289–1337. Elsevier. 10.1016/B978-0-444-59531-7.00019-3.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008. "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies* 45 (3): 693–729.

Harding, John P., Stuart S. Rosenthal, and C. F. Sirmans. 2007. "Depreciation of Housing Capital, Maintenance, and House Price Inflation: Estimates from a Repeat Sales Model." *Journal of Urban Economics* 61 (2): 193–217.

Herkenhoff, Kyle F., Lee E. Ohanian, and Edward C. Prescott. 2018. "Tarnishing the golden

and empire states: Land-use restrictions and the U.S. economic slowdown." *Journal of Monetary Economics* 93: 89–109. 10.1016/j.jmoneco.2017.11.001.

Hsieh, Chang-Tai and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *The Quarterly Journal of Economics* 124 (4): 1403–1448. 10.1162/qjec.2009.124.4.1403.

Hsieh, Chang-Tai and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics* 11 (2): 1–39. 10.1257/mac.20170388.

Johnson, Harry G. 1965. "Optimal Trade Intervention in the Presence of Domestic Distortions." In *Trade, Growth and the Balance of Payments: Essays in Honor of Gottfried Haberler*, edited by Robert E. Baldwin, pp. 3–34. Chicago: Rand McNally.

Koster, Hans R. A. 2024. "The Welfare Effects of Greenbelt Policy: Evidence from England." *The Economic Journal* 134 (657): 363–401. 10.1093/ej/uead077.

Macek, James. 2024. "Housing Regulation and Neighborhood Sorting Across the United States." https://jamesmacek.github.io/research/. Working Paper.

Marantz, Nicholas J., Christopher S. Elmendorf, and Youjin B. Kim. 2023. "Evaluating California's Accessory Dwelling Unit Reforms: Preliminary Evidence and Lessons for State Governments." Tech. rep., NYU Furman Center for Real Estate & Urban Policy. https://furmancenter.org/files/Evaluating_California%E2%80%99s_Accessory_Dwelling_Unit_Reforms_508.pdf.

Margossian, Charles C. 2019. "A review of automatic differentiation and its efficient implementation." *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 9 (4): e1305. 10.1002/widm.1305.

McDonald, John F. 2009. "Calibration of a Monocentric City Model with Mixed Land Use and Congestion." *Regional Science and Urban Economics* 39 (1): 90–96. 10.1016/j.regsciurbeco.2008.06.005.

Mei, Mike. 2022. "House Size and Household Size: The Distributional Effects of the Minimum Lot Size Regulation." https://mikemei.com/uploads/1/3/7/2/137207430/mei_mike_2022_jmp.pdf. Working Paper.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg. 2018. "Commuting, Migration, and Local Employment Elasticities." 108 (12): 3855–3890. https://pubs.aeaweb.org/doi/10.1257/aer.20151507.

Moon, Byunggeor. 2019. "The effect of FAR (floor area ratio) regulations on land values: The case of New York." *Papers in Regional Science* 98 (6): 2343–2354. 10.1111/pirs.12421.

Owens, Raymond III, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte. 2020. "Rethinking Detroit." *American Economic Journal: Economic Policy* 12 (2): 258–305.

Pacyga, Dominic A. 2009. *Chicago: A Biography*. Chicago: University of Chicago Press.

Potter, Brian and Chad Syverson. 2025. "Building Costs and House Prices." BFI Working Paper 2025-89, Becker Friedman Institute for Economics at the University of Chicago.

Redding, Stephen J. 2023. "Quantitative Urban Models: From Theory to Data." *Journal of Economic Perspectives* 37 (2): 111–138.

Redding, Stephen J. 2024a. "Quantitative Urban Economics." Working Paper 33130, National Bureau of Economic Research. https://www.nber.org/papers/w33130.

Redding, Stephen J. 2024b. "Quantitative Urban Economics." NBER Working Paper 33130, National Bureau of Economic Research.

Redding, Stephen J. and Esteban Rossi-Hansberg. 2017. "Quantitative Spatial Economics." *Annual Review of Economics* 9 (1): 21–58.

Restuccia, Diego and Richard Rogerson. 2008. "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments." *Review of Economic Dynamics* 11 (4): 707–720. 10.1016/j.red.2008.05.002.

Rossi-Hansberg, Esteban. 2004. "Optimal Urban Land Use and Zoning." *Review of Economic Dynamics* 7 (1): 69–106. 10.1016/S1094-2025(03)00056-5.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Felipe Schwartzman. 2025. "Cognitive Hubs and Spatial Redistribution." *forthcoming in American Economic Journal: Macroeconomics* 10.1257/mac.20230013.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55 (5): 999–1033. http://www.jstor.org/stable/1911259.

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *The Quarterly Journal of Economics* 125 (3): 1253–1296. 10.1162/qjec.2010.125.3.1253.

Sargent, Thomas J. and John Stachurski. 2025. "Quantitative Economics with Python using JAX." QuantEcon, online. https://jax.quantecon.org/. Jupyter Book; CC BY-SA 4.0. PDF dated Aug 04, 2025.

Sargent, Thomas J. and Ziyue Yang. 2025a. "Machine Learning a Ramsey Plan." *Quarterly Review (Federal Reserve Bank of Minneapolis)* 45 (1). 10.21034/qr.4511.

Sargent, Thomas J. and Ziyue Yang. 2025b. "Machine Learning Calvo's Optimal Plan." SSRN. 10.2139/ssrn.5183786.

Shertzer, Allison, Tate Twinam, and Randall P. Walsh. 2018. "Zoning and the Economic Geography of Cities." *Journal of Urban Economics* 105: 20–39. 10.1016/j.jue.2018.01.006.

Silva, J. M. C Santos. and Silvana Tenreyro. 2006. "The Log of Gravity." *The Review of Economics and Statistics* 88 (4): 641–658. 10.1162/rest.88.4.641.

Tan, Ya, Zhi Wang, and Qinghua Zhang. 2020. "Land-use regulation and the intensive margin of housing supply." *Journal of Urban Economics* 115: 103199. https://doi.org/10.1016/j.jue.2019.103199.

Wachter, Andreas and Lorenz T. Biegler. 2006. "On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming." *Mathematical Programming* 106 (1): 25–57.

Wang, Zi, Xingcheng Xu, Yanqing Yang, and Xiaodong Zhu. 2024. "Optimal Trade and Industrial Policies in the Global Economy: A Deep Learning Framework." Working Paper tecipa-781, University of Toronto, Department of Economics. Also available as arXiv:2407.17731.

APPENDIX A

APPENDIX

A.1 Proof that land use regulations are not first-best policies

I use Fajgelbaum and Gaubert (2020)'s framework and proof for the optimality of location-workplace specific income transfers in achieving the social planner allocation to prove that land use regulations cannot achieve the same first-best allocation as the social planner. Fajgelbaum and Gaubert first write down the problem of the social planner, who does not have prices in his problem and only needs to consider goods feasbility and spatial equilibrium. They show that for a chosen level of their proposed policy of location-workplace income transfers, the first-order conditions implied by the social planner's welfare maximization problem coincide with the equilibrium conditions of general equilibrium spatial model. Fajgelbaum and Gaubert's spaital model is very similar to mine, and their arguments are directly adaptable to mine. I write down the problem of a social planner who can chose a menu (C_{ij}, H_{ij}) of consumption at each location-workplace pair. The planner in the model still has to respect spatial equilibrium- agents in the model with idiosyncratic shocks , given the menu of (C_{ij}, H_{ij}), can choose the location that gives them the highest utility. The planner also has to respect goods feasibility.

Writing down the FOCs of the planner's problem, we compare the FOCs of the social planner's problem with the FOCs of my decentralized general equilibrium model where the policymaker can choose only land use regulations. I allow the policymaker to adjust both T_{Ci} and T_{Ri}, giving the policy maker instruments on both residential and commercial real estate. If some choice of T_{Ri}, T_{Ci} can coincide with the social planner's FOCs, then that choice of T_{Ri} and T_{Ci} can obtain 1st best outcomes. I show that there is no choice of T_{Ri} and T_{Ci} that coincide with all social planner FOCs, and thus land use regulations cannot

attain the first best.

The social planner's problem The social planner chooses an allocation across residence–workplace pairs that maximizes ex-ante expected utility. Specifically, the planner selects: (i) the mass of people living in i and working in j, N_{ij}, (ii) a menu of individual final goods and housing consumption for those agents $(C_{ij}, h_{R,ij})$, (iii) location-level stocks of residential and commercial floor space $\{H_{Ri}, H_{Ci}\}$ produced with land inputs $\{L_{Ri}, L_{Ci}\}$ and improvements $\{X_{Ri}, X_{Ci}\}$, and (iv) the aggregate utility level U. These choices must satisfy:

1. **Final-goods feasibility**, equating aggregate consumption plus construction inputs to output produced at each workplace j using effective labor $\sum_i N_{ij}$ and commercial space H_{Cj};

2. **Real-estate feasibility** at each location i :

$$H_{Ri} = Z_{Ri} L_{Ri}^{\gamma_R} X_{Ri}^{1-\gamma_R} \qquad \text{and} \qquad H_{Ci} = Z_{Ci} L_{Ci}^{\gamma_C} X_{Ci}^{1-\gamma_C},$$

3. **Population accounting** $(\sum_{i,j} N_{ij} = \bar{N})$; and

4. **Spatial-equilibrium consistency with idiosyncratic shocks**: Given individual Frechet shocks with dispersion ε, expected utility is

$$U^{\varepsilon} = \frac{1}{\bar{N}} \sum_{i,j} N_{ij} V_{ij}^{\varepsilon}, \qquad V_{ij} \equiv \frac{B_i}{d_{ij}} \left(\frac{C_{ij}}{\alpha}\right)^{\alpha} \left(\frac{h_{R,ij}}{1-\alpha}\right)^{1-\alpha}.$$

   This is equivalent to the planner choosing a choosing a level of expected utility U, subject to

$$U = \bar{N}^{1/\varepsilon} N_{ij}^{-1/\varepsilon} \frac{B_i}{d_{ij}} \left(\frac{C_{ij}}{\alpha}\right)^{\alpha} \left(\frac{h_{R,ij}}{1-\alpha}\right)^{1-\alpha} \quad , \forall i, j$$

Choosing U, $\{N_{ij}, C_{ij}, h_{R,ij}\}_{i,j}$, $\{H_{Ri}, H_{Ci}, X_{Ri}, X_{Ci}\}_i$, the social planner maximizes ex-ante expected utility subject to feasibility and spatial-equilibrium consistency. The planner's problem is thus written as:

$$
\begin{aligned}
\max \ \mathcal{L}^{FB} \ = \ & U \\
& - \sum_i \sum_j \lambda_{u,ij} \left[U - \bar{N}^{1/\varepsilon} N_{ij}^{-1/\varepsilon} \frac{B_i}{d_{ij}} \left(\frac{C_{ij}}{\alpha} \right)^{\alpha} \left(\frac{h_{R,ij}}{1-\alpha} \right)^{1-\alpha} \right] \\
& - \lambda_c \left[\sum_i \sum_j C_{ij} N_{ij} + \sum_i X_{Ci} + \sum_i X_{Ri} - \sum_j A_j \Big(\sum_i N_{ij} \Big)^{\beta} H_{Cj}^{1-\beta} \right] \\
& - \sum_i \lambda_{H_R,i} \left[\sum_j h_{R,ij} N_{ij} - Z_{Ri} \, L_{Ri}^{\gamma_R} X_{Ri}^{1-\gamma_R} \right] \\
& - \sum_i \lambda_{H_C,i} \Big[H_{Ci} - Z_{Ci} \, L_{Ci}^{\gamma_C} X_{Ci}^{1-\gamma_C} \Big] \\
& - \lambda_N \left[\sum_i \sum_j N_{ij} - \bar{N} \right]
\end{aligned}
$$

The constraint represented by the Lagrangian multiplier λ_C represents final goods feasibility- the total amount of the numeraire good consumed, $\sum_i \sum_j C_{ij} N_{ij}$, plus the total improvements used for residential and commercial construction $\sum_i X_{Ci} + \sum_i X_{Ri}$ must equate the actual amount of final goods produced by final goods firms.

The set of constraints $\{\lambda_{H_{Ci}}, \lambda_{H_{Ri}}\}$ over each location i imposes that quantities of commercial and residential real estate at each location chosen by the social planner must equal the supply of real estate; the amount of improvements used in construction X_{Ci}, X_{Ri} in each location must be enough to produce the amount of real estate chosen by the planner in each location.

FOCs of the social planner's problem We look at the FOCs of the planner with respect to location-specific improvements X_{Ci} and X_{Ri}. For each location i:

$$\frac{\partial \mathcal{L}^{FB}}{\partial X_{Ri}}: \quad 1 = \frac{\lambda_{H_{Ri}}}{\lambda_c} \frac{\partial H_{Ri}}{\partial X_{Ri}} \tag{A.1}$$

$$\frac{\partial \mathcal{L}^{FB}}{\partial X_{Ci}}: \quad 1 = \frac{\lambda_{H_{Ci}}}{\lambda_c} \frac{\partial H_{Ci}}{\partial X_{Ci}} \tag{A.2}$$

$$\tag{A.3}$$

Comparing social planner FOCs to general equilibrium conditions A general-equilibrium market allocation under a specific level of land use regulations T_{Ri}, T_{Ci} can achieve the first best only if the market allocation coincides with the social planner FOCs. Suppose, for the sake of contradiction, that there exists a level of T_{Ri}, T_{Ci} that matches the planner FOCs.

The social planner allocation implies the usual shadow-price equalization in general equilibrium: $\frac{\lambda_{H_{Ri}}}{\lambda_c} = p_{Ri}$ and $\frac{\lambda_{H_{Ri}}}{\lambda_c} = p_{Ci}$, where p_{Ri} and p_{Ci} are the prices of residential and commercial real estate. As such, in general-equilibrium allocations that attain the first best, it must be that

$$1 = p_{Ri} \frac{\partial H_{Ri}}{\partial X_{Ri}} \tag{A.4}$$

$$1 = p_{Ci} \frac{\partial H_{Ci}}{\partial X_{Ci}} \tag{A.5}$$

The general equilibrium conditions for real estate developer optimization, under land use

regulations $T_{Ri,T_{Ci}}$, are:

$$1 = p_{Ri} T_{Ri} \frac{\partial H_{Ri}}{\partial X_{Ri}} \tag{A.6}$$

$$1 = p_{Ci} T_{Ci} \frac{\partial H_{Ci}}{\partial X_{Ci}} \tag{A.7}$$

For the general equilibrium conditions and the social planner FOCs to coincide, it must be that $T_{Ri} = 1, T_{Ci} = 1$ across all locations. However, this would imply a decentralized equilibrium with no government intervention at all. Fajgelbaum and Gaubert prove that the laissez-faire equilibrium is not the first best because marginal utilities of consumption across locations are not equivalent. As such, the policy in the form of T_{Ri} and T_{Ci} cannot achieve the first-best social planner allocation that internalizes all productivity and amenity externalities.

A.2 Wedge τ is observationally equivalent to a quantity limit

I show that wedge τ_ω is observationally equivalent to a quantity limit, which is implemented in Cook County as floor-area ratio limits. I suppress the ω indices for brevity, reintroducing them only when required for clarity; all lowercase variables should be understood to be parcel-specific.

Suppose the price-taking developer faces a binding quantity constraint on real estate output: $h \leq \bar{h}$. The developer's optimization problem under the quantity limit $\bar{h}$ is:

$$\max_x \; p\,h(x,l;z) - x - q\,l \quad \text{s.t.} \quad h(x,l;z) \leq \bar{h}, \quad \mu \geq 0.$$

Under the Lagrangian $\mathcal{L} = p\,h - x - q\,l + \mu(\bar{h} - h)$, with Lagrangian multiplier $\mu > 0$,

$$1 = (p-\mu)\,\frac{\partial h}{\partial x} = (p-\mu)\,\gamma\,\frac{\bar{h}}{x} \quad \Rightarrow \quad x^{\text{limit}} = \gamma\,(p-\mu)\,\bar{h}.$$

, where I denote x^{limit} as the developer's choice under a binding quantity limit. Land price still drives profits to zero; $ql = ph - x$. Under the quantity limit land rent is:

$$q^{\text{limit}}\,l = p\,\bar{h} - x^{\text{limit}} = p\,\bar{h} - \gamma\,(p-\mu)\,\bar{h} = \big[(1-\gamma)\,p + \gamma\,\mu\big]\bar{h}.$$

Now we look at the problem of a developer who instead faces a price wedge τ:

$$\max_x \; \tau\,p\,h(x,l;z) - x - q\,l, \qquad 1 = \tau\,p\,\frac{\partial h}{\partial x} = \tau\,p\,\gamma\,\frac{h}{x} \Rightarrow x^{\text{wedge}} = \gamma\,\tau\,p\,h.$$

, where x^{wedge} is the developer's choice if he faces a virtual distortion in the form of a price wedge. Under the wedge formulation land rents q^{wedge} are:

$$q^{\text{wedge}}\,l = p\,h - x^{\text{wedge}} = (1-\gamma\,\tau)\,p\,h.$$

I now show that $x^{\text{limit}} = x^{\text{wedge}}$ and $q^{\text{limit}} l = q^{\text{wedge}} l$ when we set $\tau = \frac{p-\mu}{p}$; the wedge is observationally equivalent to the quantity limit. Substituting in $\tau = \frac{p-\mu}{p}$ into our expression for $x^{\text{wedge}}, q^{\text{wedge}}$:

$$x^{\text{wedge}} = \gamma \, \tau \, p \, \bar{h} = \gamma \, (p - \mu) \, \bar{h} = x^{\text{limit}}$$

$$q^{\text{wedge}} \, l = (1 - \gamma \, \tau) \, p \, \bar{h} = \big[(1-\gamma)p + \gamma\mu\big] \bar{h} = q^{\text{limit}} l$$

The quantity limit is observationally equivalent to a price wedge: it delivers the same x, the same $h = \bar{h}$, and the same land values. The distortion τ can reflect both the housing quantity and land-value effects of quantity constraints imposed by different types of land use regulations that are difficult to observe in reality. τ admits two interpretations: (i) a "virtual" distortion in the developer's problem that rationalizes the observed construction and land values, or (ii) a per–unit tax $(1 - \tau)$ on developers redistributed to landowners .

A.3 Wedge τ corresponds to a shadow price of extra housing in the presence of a quantity limit

Suppose that there is a FAR constraint on an otherwise unrestricted builder: $\frac{H}{l} \leq F\bar{A}R$. The FAR constraint is an explicit quantity constraint. The builder takes real estate prices p as given and cannot influence price. The FAR-constrained builder's problem can be written as:

$$\max_x \mathcal{L} = pH(l,x) - x - \lambda(\frac{H(l,x)}{l} - F\bar{A}R)$$

Which gives the FOC wrt x:

$$\left(p - \frac{\lambda}{l}\right)\frac{\partial H}{\partial x} = 1$$

Compare this to the builder's FOC when he faces wedge τ:

$$\tau p \frac{\partial H}{\partial x} = 1$$

When $\tau = p - \frac{\lambda}{l}$, the same choice variables H, x are obtained from both the quantity-constrained producer and the wedge-constrained producer. We see that a higher value of λ, indicating a higher value of relaxing the quantity constraint, corresponds to a lower level of τ, which corresponds to higher stringency.

A.4 Deriving resident and firm market access from commuting probabilities

This appendix derives the resident market access (RMA) and firm market access (FMA) terms from the residence-workplace choice probabilities in the model.

Starting from equation (3.3), the probability that an individual chooses to live in location i and work in location j is

$$\pi_{ij} = \frac{\left(\frac{B_i e_j}{d_{ij} p_{Ri}^{1-\alpha}}\right)^{\epsilon}}{\sum_{m,n \in I^2} \left(\frac{B_m e_n}{d_{mn} p_{Rm}^{1-\alpha}}\right)^{\epsilon}}, \tag{A.8}$$

where $e_j \equiv w_j + \bar{y}$ is income for an individual working in location j.

For compactness, define the common denominator

$$\Phi \equiv \sum_{m,n \in I^2} \left(\frac{B_m e_n}{d_{mn} p_{Rm}^{1-\alpha}}\right)^{\epsilon}. \tag{A.9}$$

Using (A.8), residential population in location i can be written as

$$R_i = \bar{N} \sum_j \pi_{ij} \tag{A.10}$$

$$= \bar{N} \sum_j \frac{\left(\frac{B_i e_j}{d_{ij} p_{Ri}^{1-\alpha}}\right)^{\epsilon}}{\Phi} \tag{A.11}$$

$$= \bar{N} \frac{\left(\frac{B_i}{p_{Ri}^{1-\alpha}}\right)^{\epsilon} \sum_j \left(\frac{e_j}{d_{ij}}\right)^{\epsilon}}{\Phi}. \tag{A.12}$$

This motivates the definition of *resident market access* at residential location i:

$$RMA_i \equiv \sum_j \left(\frac{e_j}{d_{ij}} \right)^{\epsilon}. \tag{A.13}$$

Substituting (A.13) into (A.12) yields

$$R_i = \bar{N} \frac{\left(\frac{B_i}{p_{Ri}^{1-\alpha}} \right)^{\epsilon} RMA_i}{\Phi}. \tag{A.14}$$

Resident market access therefore summarizes the commuting-cost-discounted access of residents in location i to high-income workplace locations. Holding local amenities and housing prices fixed, a location with higher RMA_i attracts more residents because it offers better access to earnings opportunities elsewhere in the city.

A parallel derivation yields firm market access. Workplace employment in location j is

$$N_j = \bar{N} \sum_i \pi_{ij} \tag{A.15}$$

$$= \bar{N} \sum_i \frac{\left(\frac{B_i e_j}{d_{ij} p_{Ri}^{1-\alpha}} \right)^{\epsilon}}{\Phi} \tag{A.16}$$

$$= \bar{N} \frac{e_j^{\epsilon} \sum_i \left(\frac{B_i}{d_{ij} p_{Ri}^{1-\alpha}} \right)^{\epsilon}}{\Phi}. \tag{A.17}$$

This motivates the definition of *firm market access* at workplace location j:

$$FMA_j \equiv \sum_i \left(\frac{B_i}{d_{ij} p_{Ri}^{1-\alpha}} \right)^{\epsilon}. \tag{A.18}$$

Substituting (A.18) into (A.17) gives

$$N_j = \bar{N} \frac{e_j^{\epsilon} FMA_j}{\Phi}. \tag{A.19}$$

Firm market access therefore summarizes the commuting-cost-discounted access of firms in location j to residential locations that are attractive to workers. Locations with high amenities B_i and low housing prices p_{Ri} contribute more to FMA_j, because they make it easier for firms in j to draw labor supply from surrounding residential locations.

The two market access objects summarize the spatial linkages embedded in the commuting system. RMA_i captures how well a residential location is connected to income-generating workplaces, while FMA_j captures how well a workplace is connected to worker-attractive residential locations. These summaries are useful for interpreting the spatial pattern of optimal land use regulations in the model.

A.5 Automatic differentiation

Automatic differentiation (AD) is a method for evaluating derivatives by decomposing a function into a sequence of elementary operations and systematically applying the chain rule to that sequence. It differs from symbolic differentiation, which manipulates algebraic expressions, and from numerical differentiation, which approximates derivatives through finite differences. Because AD operates directly on the sequence of computations used to evaluate the function, it produces derivatives that are exact up to machine precision while avoiding the truncation error and repeated function evaluations associated with finite-difference methods.

In my setting, AD is useful because the equilibrium mapping $F(X)$ is composed of many nested nonlinear and multiplicative expressions. AD represents this mapping as a graph: starting from the input vector X, the program constructs intermediate variables through operations such as addition, multiplication, exponentiation, and trigonometric transformations, and then maps those intermediates into the final output. Derivatives are then obtained by propagating local derivative information through the graph via the chain rule.

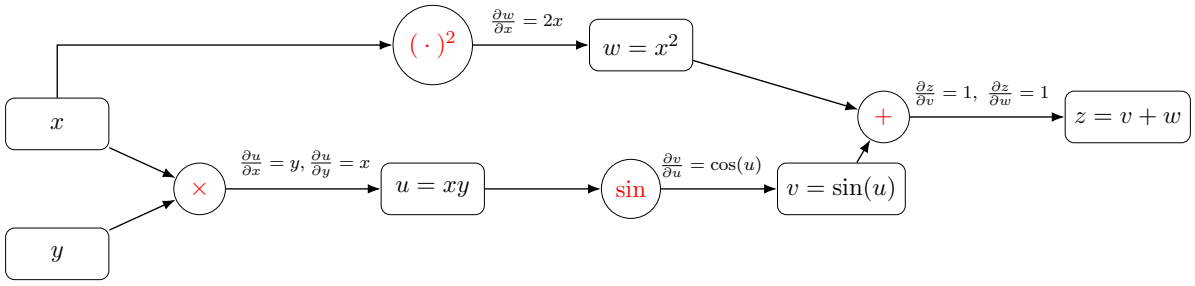

Figure A.1: Illustration of a computational graph for $f(x, y) = \sin(xy) + x^2$

Figure A.1 illustrates these ideas for the simple function $f(x, y) = \sin(xy) + x^2$. Defining intermediate variables $u = xy$, $v = \sin(u)$, $w = x^2$, and $z = v + w$, the derivative of z is obtained by combining the local derivatives associated with each edge of the graph. More generally, AD automates this same procedure for large nonlinear systems, making the

computation of gradients and Jacobians feasible in high-dimensional problems such as the optimal policy problem studied here.

AD can be implemented in two principal modes. In forward mode, the derivative is carried alongside the original calculation: as each intermediate variable is computed, the program also updates how a small change in an input would affect that intermediate and, in turn, the final output. In reverse mode, the function is first evaluated in the usual forward direction, after which the derivative calculation proceeds backward through the computational graph, accumulating how each intermediate variable contributes to the derivative of the final output. Forward mode is typically advantageous when the number of inputs is limited, whereas reverse mode is particularly useful when differentiating a scalar-valued objective with respect to many inputs.

Reverse mode is especially widely used in modern machine learning because training typically requires the gradient of a single scalar loss function with respect to a very large number of parameters. In machine learning settings, reverse-mode AD is commonly referred to as backpropagation. Modern AD packages typically implement both forward and reverse modes. The CasADi package can automatically choose between forward- and reverse-mode directional derivatives.

A.6 Implementation details for AD-based optimization

A.6.1 Quasi-Newtonian approximation of the Hessian

While AD can in principle also provide second derivatives, forming and factorizing the full Hessian at the problem's scale adds to computational costs. Instead, I use a quasi-Newtonian Broyden–Fletcher–Goldfarb–Shanno (BFGS) method to approximate second-order curvature. This method is an application of the multivariate mean-value theorem to iteratively approximate the Hessian of the objective from changes in the gradient across search iterations

.

Let T_R^k be the policy vector at iteration k. After updating to T_R^{k+1}, define the step size $s_k = T_R^{k+1} - T_R^k$ and the gradient change $y_k = \nabla_{T_R}\mathcal{O}(T_R^{k+1}) - \nabla_{T_R}\mathcal{O}(T_R^k)$. By the multivariate mean–value theorem, along the line segment from T_R^k to T_R^{k+1} there exists $\bar{t} \in (0,1)$ such that

$$\nabla_{T_R}\mathcal{O}(T_R^{k+1}) - \nabla_{T_R}\mathcal{O}(T_R^k) \;=\; \nabla^2_{T_R}\mathcal{O}\big(T_R^k + \bar{t}\, s_k\big)\, s_k.$$

The Hessian approximation maps the step $s_k = T_R^{k+1} - T_R^k$ to the observed gradient change $y_k = \nabla_{T_R}\mathcal{O}(T_R^{k+1}) - \nabla_{T_R}\mathcal{O}(T_R^k)$, using the Hessian $\nabla^2_{T_R}\mathcal{O}\big(T_R^k + \bar{t}\, s_k\big)$ somewhere along step. This motivates the secant condition which the BFGS approximation to the Hessian, $H_k \approx \nabla^2_{T_R}\mathcal{O}(T_R^k)$ must satisfy:

$$H_{k+1}\, s_k \;=\; y_k.$$

In one dimension, the secant condition uniquely determines the approximation of the second derivative. In many dimensions, H_{k+1} is an $n \times n$ symmetric matrix with $n(n+1)/2$ unknown entries, while $H_{k+1}s_k = y_k$ supplies only n equations. Infinitely many matrices satisfy the secant condition. The BFGS algorithm selects the Hessian approximation that

enforces the secant condition exactly, preserves symmetry and positive definiteness; and changes H_k outside the subspace informed by the new data (s_k, y_k) as little as possible. BFGS show that the directional update that satisfies these conditions is:

$$H_{k+1} - H_k = \frac{y_k\, y_k^\top}{y_k^\top s_k} - \frac{H_k\, s_k\, s_k^\top H_k}{s_k^\top H_k\, s_k}$$

The first term updates the Hessian approximation given the step s_k, and the change in the gradient y_k, using the secant condition to map s_k to y_k. The second term removes the old information in H_k previously implied along step direction s_k . The updated Hessian approximation H_{k+1} remains positive definite, so H_{k+1}^{-1} provides a second-order transformation of the policy gradient for the update of our policy vector given the gradient of the objective with respect to the policy vector $\nabla_{T_R}\mathcal{O}(T_R^{k+1})$.

A.6.2 Optimization under constraints on T_{Ri}

In solving for the optimal regulation, I impose the constraint $0 \leq T_{Ri} \leq 1$ for each location i. The upper bound $T_{Ri} = 1$ corresponds to no regulation at i, while $T_{Ri} = 0$ corresponds to a ban on construction. To ensure that the numerical solution respects these bounds, I switch from the unconstrained Newtonian optimizer to the Interior Point Optimization (IPOPT) solver implemented by Wachter and Biegler (2006) when a candidate policy is within 0.02 of the bounds. IPOPT is a computational constrained-optimization solver that handles bound constraints via barrier methods. The IPOPT algorithm converts constrained optimization problems into a nonconstrained problem with barrier terms that penalize proximity to the bounds, thereby keeping each T_{Ri} within $(0, 1)$.

Specifically, the interior-point approach solves a sequence of "barrier subproblems" that include logarithmic penalties for approaching the bounds. Within each iteration of the outer loop where unconstrained optimization violated the bounds on T_{Ri}, the IPOPT package solves a series of barrier subproblems for a series of barrier parameters μ that tend towards zero to obtain a new search direction. Given barrier parameter $\mu > 0$, the solver minimizes[1] a modified objective:

$$\min_{T_R} -\mathcal{O}(T_{Ri}) \; - \; \mu \sum_i \Big[\ln\big(T_{Ri} - 0\big) + \ln\big(1 - T_{Ri}\big) \Big]$$

where $\mathcal{O}(T_{Ri}) = O(X(T_R), T_R)$ denotes the original (government expected utility) objective function. Due to the properties of the log function, the added terms $-\mu \ln(T_{Ri})$ and $-\mu \ln(1 - T_{Ri})$ tend to infinity when any T_{Ri} approaches its lower bound 0 or upper bound 1. For any fixed $\mu > 0$, the solution of this barrier problem is forced to satisfy $0 < T_{Ri} < 1$ strictly for all i by the logarithmic boundary penalty.

1. Given that the IPOPT is a minimizer, I minimize the negative of the objective function.

The IPOPT algorithm uses the T_R that solves each boundary subproblem with boundary parameter μ as the initial guess for the next boundary subproblem with boundary parameter μ' that is strictly smaller than the previous μ. As μ goes toward zero across iterations of the boundary subproblem, the solution converges to the original unconstrained problem of $\min_{T_R} -\mathcal{O}(T_{Ri})$ that satisfies the non-penalized optimality conditions and yields T_R values within the $[0, 1]$ bounds.

A.7 Optimal policies under $\phi_{amen} = 0.15, \phi_{prod} = 0.07$

In Figure 4.1, all optimal policies are plotted on the same color scale to compare optimal policy across different scenarios. The optimal policy for $\phi_{amen} = 0.15, \phi_{prod} = 0.07$ shown in panel (B) has relatively low contrast. Below, I plot a graph with a color scale centered around the mean of optimal policies for this specific case.

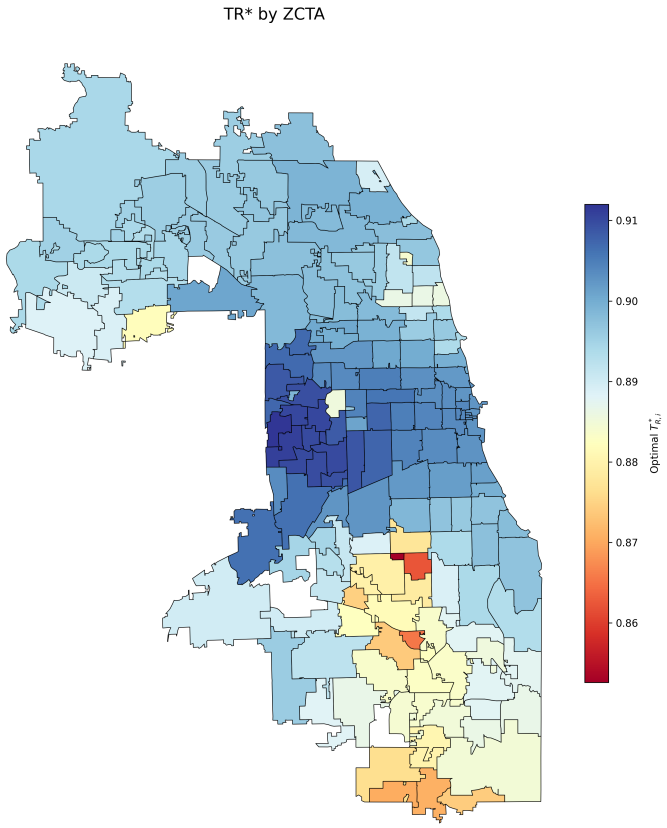

Figure A.2: Optimal land use regulations T_{Ri} across space, $\phi_{amen} = 0.15, \phi_{prod} = 0.07$