Essays on Housing and Demography

by

Mike Mei

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2023

Doctoral Committee:

Professor Hoyt Bleakley, Chair
Assistant Research Scientist Gabriel Ehrlich
Professor Jonathan Levine
Assistant Professor Ana Reynoso

Mike Mei
mikemei@umich.edu
ORCID iD: 0000-0002-2887-7715

DEDICATION

To my life partner Yichen Zhang

ACKNOWLEDGMENTS

There are many who have helped me to get where I am today with this research. The biggest contributor has been my committee, chaired by Hoyt Bleakley. I have benefited immensely from Hoyt's depth and breadth of economic insight (ranging across different subfields of urban, history, labor, development and others). He encouraged me to do research in housing and real estate because he thought it was an understudied topic. His genuine interest and belief in my research really propelled me to do well. His suggestion that I start sharing my research at conferences really paid off.

Ana Reynoso's presence on the committee was invaluable. She was one of the first people to really believe in the promise of my papers even though my initial research ideas were vague. Her advice about starting with simple models and then building up to complex models really had an impact on me. She was really great at pushing me to be clear in my presentations and my writing. Hearing her perspective on working as junior faculty was also incredibly helpful.

Committee member Gabriel Ehrlich was somebody I could really count on. He was crucial in helping me steer my research in a clear direction after getting some initial discouraging feedback as a graduate student. His attention to detail, which often involved discussions about footnotes and minute details related to specific parameters of the model, kept me accountable to rigorous standards. He really cared a lot about my progress and my job market outcome.

Committee member Jonathan Levine's influence on my research has been substantial. My study of Houston in my first chapter and my study of vacancy chains in my third chapter came out of conversations with him about land use regulations and existing holes in the urban literature. His understanding of land use regulations and their nuances offered a fresh perspective that greatly complemented the economics field.

Other than my committee, I would like to acknowledge many people who made it possible to attend the Michigan Economics Ph.D. program and finish my dissertation. It all starts with my family. My paternal grandparents immigrated to the United States in the 1980's working long hours in grueling restaurant and service jobs so that the next generation could have a better life. My maternal grandmother never learned to read or write in any language. My parents worked their way through graduate school in the United States while raising me and my sister. It is often said that we stand on the shoulders of giants, but sometimes those giants are simply the ordinary people in your life who have sacrificed so much for you.

Professor Raaj Sah at the University of Chicago was my mentor throughout my undergraduate years. He was instrumental in getting me interested in economic research. He was an incredibly warm and kind person, and it was a pleasure to have so many dinners at his place. I thoroughly enjoyed learning about economic theory by proofreading his papers. It meant so much to me that he took the feedback of an undergraduate so seriously for his research.

Professor Erik Hurst at the University of Chicago taught a Ph.D. level class about how to use microdata to answer big questions; this class continues to have an influence on my work today. He strongly supports his students and goes the extra mile to help them succeed. It was his advice on graduate school that led me to choose the Michigan Economics program over other better funded alternatives.

Anna Paulson and Gene Amromin were the economists at the Federal Reserve that I worked with very closely. They offered me an incredible amount of support as a research assistant there. It was working with these economists that I learned how to effectively work with data, especially mortgage and other consumer finance data which forms a big part of this dissertation.

John Bound offered me key mentorship and advice at crucial junctions of my graduate studies. His role as the economics PSC coordinator and advisor guided me toward a soft landing for my research. It was a pleasure to also help teach his migration economics class and to be nudged towards seriously considering a career at a liberal arts college. He was one of the key people to have believed in me in throughout my graduate studies.

Getting through the first two years of graduate school was made a lot easier through the camaraderie and mutual support of fellow graduate students in my cohort. I fondly remember the days of studying for examinations with Olga Fetisova, Jamie Fogel, Giacomo Meille, Bernardo Modenesi, Xuan Wang, and many others. Subsequent graduate student-led reading groups and research presentation groups really made it feel like we had an exciting intellectual community here that went beyond the usual seminar groups. Finally, the 2022-2023 Michigan job market candidate cohort this year offered me a lot of support, and it made that whole process feel a lot less lonely.

My work deals with deep and sometimes unsettling questions of welfare and justice. Many of these questions arise whenever protest, pandemic, or war happens in the world. The fact that the world was a tumultuous place throughout my years as a graduate student inevitably influenced my productivity and well-being. In the midst of this tumultuous world, I found the courage to push on. In particular, it was the bravery of the people of Hong Kong and the people of Ukraine that inspired me to persevere even if I felt like the odds were overwhelming.

Being able to study and do research for a Ph.D. is a privilege that is made possible by certain material conditions. For me, these material conditions came from the fruits of hard work and sacrifice by my life partner Yichen Zhang, who moved to Michigan for me and provided so much for our household. The second source comes from the work of fellow graduate students at the

Graduate Employees' Organization and at Rackham Student Government who are still working hard to improve the conditions of scholarship and labor at the University of Michigan.

This research was supported in part by a National Institute of Child Health and Human Development (NICHD) training grant through the Population Studies Center (Institute for Social Research) at the University of Michigan. The content of this dissertation is solely the responsibility of the author and does not necessarily represent the views of the NICHD.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

LIST OF APPENDICES

LIST OF ACRONYMS

2SLS Two-stage Least Squares

CPS Current Population Survey

FE Fixed effects

FHA Federal Housing Administration

HH Household

HMDA Home Mortgage Disclosure Act

IPUMS Integrated Public Use Microdata Series

IV Instrumental variables

MB Marginal benefit

MC Marginal cost

MLS Minimum lot size

MSA Metropolitan Statistical Area

OLS Ordinary Least Squares

SQFT Square feet

VA Veterans Administration

WRLURI Wharton Residential Land Use Regulation Index

WWII World War II

ABSTRACT

This dissertation is an analysis of policy issues and historical patterns in housing markets and their connections to the demographic attributes of households. Chapter 1 studies Houston's significant reduction of the minimum lot size in 1999. Using synthetic control methods, I show Houston's reduction of the MLS in 1999 led to a 12% decrease in the size of new housing and a 14% increase in the marginal cost of house size. To quantify the distributional welfare effects stemming from these incentives, I build a quantitative model with housing and demographics and show that the observed price changes induced by reductions of the MLS disproportionately help lower income and smaller households. Specifically, I find that the bottom decile of households (in terms of household size and income) gain about $25,000 more (in 2010 dollars) than the top decile from a reduction in the MLS. Finally, I show that the model's predicted locational selection of households by household size and income is consistent with empirical observations in Houston before and after the change in regulation. Chapter 2 studies the demographic determinants of housing regulations. I show that places in the United States which had larger fertility booms developed more housing regulations. Using WWII-induced variation in the size of the Baby Boom, I show that this relationship is plausibly causal and that the empirical patterns are consistent with political economy models of housing regulation that are modified to include fertility booms. Chapter 3 analyzes the distributional effect of new housing construction by tracing housing vacancy chains. I construct matched buyers/sellers chains, all matched with household specific characteristics reported from mortgage application data. I find that the interlink income elasticities on the chain are relatively low, but slightly higher than those estimated using census tract level data. The evidence is still consistent with the idea that new housing construction can have relatively immediate vacancy effects for lower socioeconomic status groups.

CHAPTER 1

House Size and Household Size: The Distributional Effects of the Minimum Lot Size Regulation

1.1 Introduction

One of the most common residential land use regulations in the United States is the minimum lot size (MLS). For single family housing, the MLS sets a minimum amount of land that each unit of housing must be built upon. Given the demand for single family housing, this regulation has visible effects on the character of housing and use of space where the MLS may be binding. The MLS has two direct effects: first, it bundles additional land with each house that may otherwise not have been acquired, which increases housing prices directly through the requirement for builders to acquire more land. The second direct effect is the increase in the size of housing demanded since it is cheaper on a marginal basis to build larger, as the land needed to build larger would have been already acquired under a large MLS regulation. Thus, perhaps counterintuitively, the marginal price of an additional unit of house size (say, square feet) decreases under a (sufficiently large) MLS regulation. Hence, an MLS regulation can increase demand for housing size by distorting the marginal price of size to be lower than it would otherwise be, causing substitution effects into larger housing. While the first channel has been studied in terms of effects on overall housing prices, there has been less work on its heterogenous effects. The second channel has been largely overlooked.

These two channels of the MLS regulation are important to understand primarily because they have substantial implications for the heterogeneity in welfare losses. Specifically, households differ in income, family size, and age. To the extent that wealthier, larger, and older households already demand larger houses, the effects of these regulations on their welfare should be substantially smaller than households who are poorer, smaller, or younger. Thus, the characteristics of the family serve as shifters of demand for housing size, which can amplify welfare losses. The model in this paper suggests that a minimum lot size deregulation can differentially affect some households, in

2010 dollar terms, about $25,000 more than other households. Taking into account changes in asset prices for homeowners, the heterogeneity is substantially larger.

This paper studying a specific housing regulation and the effect on housing size is partly motivated by historical trends and perspectives. As seen in Figure 1, between 1950 and 2010, the physical size of houses in the United States increased substantially.[1] At the same time, average household size (number of people in the household) has been falling. Housing regulations, specifically MLS regulations, may have played a role in changing the various cost margins of housing supply and incentivizing larger houses. Because the households' optimal choice of house size is dependent on household size and income, there is an open question about the welfare costs (and their heterogeneity) due to these MLS regulations, as well as the channels in which the characteristics of the family can amplify these costs.

Figure 1.1: Household Size and Size of New Housing

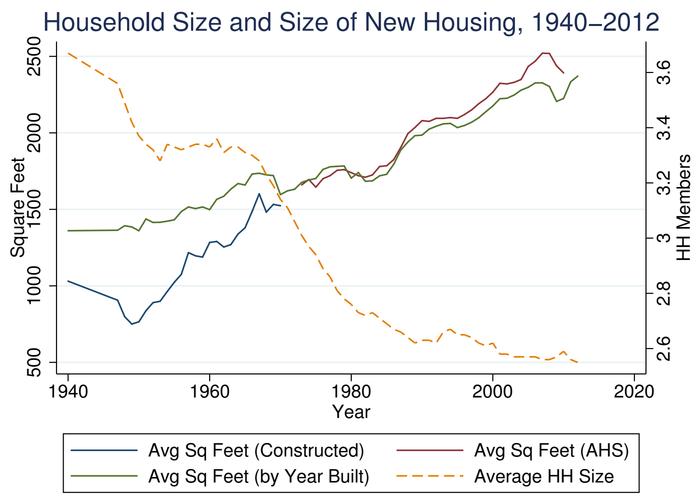

Notes: The AHS measure is the average square feet of new housing starts as measured by American Housing Survey from the U.S. Census Bureau. The Constructed Historical Estimate is calculated from data in the report titled "Historical Statistics of the United States, Colonial Times to 1970" published by the U.S. Census Bureau. The measure is constructed by taking the total square feet associated with construction contracts that year, divided by the number of housing starts. The total square feet measure is adjusted for missing states using overlapping years from the NBER source data "Source Book of Statistics Relating to Construction" by Lipsey and Preston (1966). Finally, the contracted square feet is also adjusted for the portion attributable to additions or renovations rather than new housing; this is done by multiplying by the percentage of value of those contracts that are attributable to new housing only. Corelogic supplies property tax and characteristic data from 2015-2016. By analyzing their characteristics and when they were built, we see that of the houses that survive until this day, by "year built", there has been a similar steady increase in average housing size. That is, the older housing today amongst the housing stock today is physically smaller.

In addition, there are many real-world concerns that motivate the analysis of these issues. Younger families across the United States have reported difficulty finding smaller starter homes. Their reported demographic characteristics–with lower family size (Figure 1.1)–suggests that these generations may desire smaller housing. Also, the rise in house size exacerbates environmental externalities that are already well-known: that is, larger houses require more energy to heat and cool, and require more land area and natural resources to build. In addition, to the extent that these

[1]This rise is not adequately explained by demographic covariates like income and household size (see Appendix A.2).

housing regulations discourage higher density and make it harder to live closer to central business districts, there are environmental and congestion externalities involved with transportation.

In this paper, I take a structural approach to understanding the heterogeneity in welfare losses due to the MLS regulation. This is in comparison to existing reduced form methods in the literature, like hedonic regressions to identify effects on prices. A structural approach allows for a direct simulation of a policy counterfactual, which is used to calculate welfare effects across the income and demographic distribution. In contrast, hedonic regressions compare similar houses, and merely identifies the overall change in housing prices conditional on housing characteristics. The hedonic regression approach fails to account for the long run re-optimization of the household in terms of house size. Simply put, the relevant welfare calculation should not be comparisons of similar houses, but of the same households who may change their optimal choice of housing size due to the policy. If the traditional hedonic regression model is used to measure welfare loss, that model therefore tends to overestimate the impact of MLS regulations because it does not account for households re-optimizing. Furthermore, the traditional hedonic model does not directly account for the main determinants (household size, income, age) of household heterogeneity in terms of their re-optimization decision.

The model used in this paper incorporates a standard household lifecycle model, and demographic (household size) and income changes through the changing characteristics of overlapping generations. The model departs from traditional preference structures (like CES) in order to capture many important features that are relevant for quantifying the heterogeneity in welfare costs: the nonlinear Engel curves for house size, age-dependent housing demand, and shifters (either because of preferences or technology) of house size demand over time. In addition, the estimation and simulations both depend on cross-sectional data relationships from the Census, as well as estimates of the impacts of the MLS on the marginal cost of an additional unit of housing size using reduced form estimates from a natural experiment in Houston.

The Houston natural experiment is a key part of validating the model and learning what experiment to simulate with the model. The Houston reduced form analysis estimates house size demand elasticities by looking at housing size characteristics in Houston before and after a reduction in the MLS in 1998, relative to a synthetically created control city generated using standard synthetic control methods. Houston serves as a suitable natural experiment for a variety of reasons. First, Houston has no traditional zoning and relatively few housing regulations (even though many restrictions remain in place due to private convenants and other regulations). Thus, it is a setting in which a relaxation of the MLS regulation may have an observable effect since other zoning regulations which would otherwise affect house size are not present in that jurisdiction. Second, Houston's MLS regulation was reduced from a sizable 5000 square feet down to as low as 1400 square feet, a reduction that had significant positive effects on the quantity of smaller lots that were developed in the subsequent periods. Using a variety of difference-in-difference and synthetic control methods, I find that the 1998 deregulation significantly decreased the size of new housing built in Houston by about 11% and increased the marginal cost of house size by about 14%.

The layout of this paper is as follows: The historical context and literature about this topic is discussed. The theoretical mechanisms of the minimum lot size on housing size are detailed in a quantitative model of the housing market, which motivates the study of how housing size demand reacts to changes in prices. The simulation inputs into the model are disciplined by an event study analysis of the Houston minimum lot size deregulation in 1999, as well as cross-sectional Census data. I show the simulated welfare results on households and their distributional consequences. Finally, I analyze a natural prediction due to the substantial heterogeneity, which is the selection of demographic variables before and after the regulation change, relative to comparison jurisdictions. In other words, in line with the model's predictions, I find that Houston's households are smaller and less wealthy than they otherwise would be, and that these effects come from selection into and out of Houston.

1.1.1 Background and Literature

1.1.1.1 U.S. Housing Regulations

Housing regulations have been a large topic in the urban economics and urban planning literature, with Glaeser and Gyourko (2018), Gyourko et al. (2008), Albouy and Ehrlich (2018), and Ganong and Shoag (2017) covering important measures and costs of housing regulations. Hirt (2015) provides a more detailed historical perspective of US zoning laws. The precise impact of specific regulations like minimum lot sizes is studied in Zabel and Dalton (2011) and Gray and Millsap (2020). However, the literature has not adequately covered the precise impact of housing regulations on house size demand nor the relevant marginal prices for house size demand, nor the implications of house size demand across demographics.

Existing research has theorized that demographic trends have important impacts on the housing market. Mankiw and Weil (1989) predicted a housing bust after the Baby Boom. Banks et al. (2015) theorizes demographic shifters of housing consumption across the lifecycle. The demographic context coming out of the 1960's is the end of the Baby Boom period, a period roughly between 1940 and 1965 where U.S. fertility rates broke its long term declining trend and started increasing. This boom reversed course by the 1970's, when U.S. fertility rates were declining again. I do not take a particular stance on the underlying reason behind changes in family size.[2]

[2] Many economic explanations have been provided for the cause of the Baby Boom, including the delay of fertility decisions from World War II [Doepke et al. (2015)], technological innovations in the household [Greenwood et al. (2005)], and maternal health innovations [Albanesi and Olivetti (2016)].

As is true in many of these papers, I take family size to be an exogenous shifter. The 1960's was also a time of increased mass suburbanization and additional population pressures in the form of mass movement away from city centers into larger suburban homes, which was likely accelerated by technological innovations like the widespread adoption of the automobile and the construction of the interstate highway system as noted in Fischel (2004). With the development of these new communities came concerns about the future trajectory of neighborhoods. It was shortly after the 1960's that many of the new housing regulations we see today formed. Economic historians see many different reasons for this change. Attitudes began to change against local growth, possibly because of a new realization that growth could depress housing values. Fischel (2004) notes that it was a combination of environmental concerns and uneasiness about racial diversity that motivated communities to start to severely restrict development. What ties these explanations together is that these concerns may have ultimately been induced by the economic and political environment created locally by demographic changes; specifically, communities that had large single family houses and open green spaces had plenty of incentives to keep their neighborhoods that way.

Changes in legal thought also spurred restrictive regulations on the construction of housing. Both Ganong and Shoag (2017) and Fischel (2004) write that the Mount Laurel decisions (1975 and 1985) in the Supreme Court of New Jersey were symbolic of a regulatory environment in which courts often were only hostile to regulations that were obviously or intentionally exclusionary; broad housing regulations like minimum lot sizes and open space requirements became legally accepted. In summary, the Mount Laurel decisions were ones where the plaintiffs won a small battle to build smaller affordable housing, but in doing so, unintentionally created the incentives for many communities to pass even broader regulations that circumvented the limited legal restrictions on housing regulations. Hence, roughly speaking, the time series of increasing housing regulations matches the time series of increasing house size.

1.1.1.2 Housing Size

The history of increasingly larger homes in the United States is a rather complicated one. Hirt (2015) write that part of the demand for larger homes come from deeply embedded preferences that are core to the notion of a wild American frontier (which stands in contrast to European cities). This is in line with empirical research that supports conspicuous consumption models of residential homes, as in Bellet (2019).

However, much of the historical literature has discussed government policy as a cause of larger homes in cities. This includes the multiple determinants of suburbanization, as discussed in Mieszkowski and Mills (1993). More recent literature looks specifically at zoning policies, as in Schuetz (2009), who finds that restrictive zoning policies likely decreased quantity of smaller rental housing built, but the overall effects on aggregate rent levels are unclear.

1.1.1.3 Contribution to Literature

The most direct contribution I make to the literature is to analyze an overlooked aspect of housing, which is its physical size (and the changing price of housing size). This channel is important because hedonic models which merely control for housing size in price regressions ignore both the time varying aspect of the housing size coefficient with respect to a policy change, and perhaps more importantly, they ignore the distributional impacts of policy changes across the various factors that shift demand for housing size. The standard hedonic model therefore ignores the endogenous price of housing size, as well as the structural elements of how households make decisions regarding housing size. This paper is a merging of the demographic housing literature (for which we know that demographic factors can have large effects on housing size demand) and the housing regulation literature (for which we know that regulations can affect housing prices and welfare losses from these regulations can be large).

There is also substantial work in making sure the welfare calculations are realistic. With respect to heterogeneity in the minimum lot size's welfare costs over the income distribution, one of the most relevant factors is the shape of the Engel curve for house size. In Appendix A.2.1 I show that housing is a strong necessity good (which is a good where demand increases per unit of income, but at a diminishing rate), a fact which is also consistent with existing work like Albouy et al. (2016). Hence, standard models that use homothetic preferences to model housing ignore this widely known aspect of the housing Engel curve; the model in this paper allows for housing size to be a necessity good, the extent to which the Engel curves are not linear will be disciplined by the data. I document that the use of nonhomothetic preferences in modeling housing consumption is a crucial aspect of simulating the relevant income effects across the income distribution, and more details about the importance of this feature is detailed in Section 1.4.1.1 about Model Fit.

1.2 Theoretical Framework

1.2.1 Basic Mechanism

The precise mechanism for how a minimum lot size regulation affects housing prices is illustrated here. Suppose a builder is thinking of building an average 2500 square foot house. Under a binding 5000 square foot minimum lot size, ϵ deviations from that 2500 square foot house represent differences in the price of labor and materials to build that home. However, under a nonbinding 1400 square foot minimum lot size, ϵ deviations from that 2500 square feet represent differences in the price of labor and materials, as well the cost of land if house size enters the utility function as its own good. Hence, within the support of a nonbinding minimum house size region, the marginal price of an additional unit of house size is higher relative to one in a regime with a binding (say,

Figure 1.2: Cost Structure Under Two Different Minimum Lot Sizes

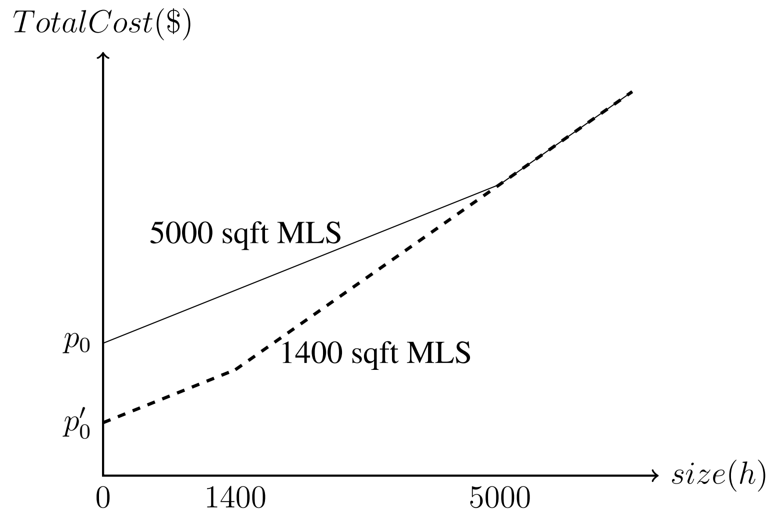

5000 square feet) minimum lot size, simply because more land needs to complement a larger house if it has not already been acquired.

The focus of this analysis is on the main channels. This analysis abstracts away some details like the requirements for open space, the ability to build vertically as opposed to horizontally, and the general equilibrium effects of the policy change on land prices. The minimum lot size directly increases the cost to build the house of the smallest base size (for example, a small studio cabin) because it requires a certain minimum amount of land to be bought. As shown in Figure 1.2, the cost p to build a house of size h can be expressed in reduced form as $p = p_0 + p_h h$ where p_0 is the base cost of a house, and p_h is the marginal cost of house size. Under a binding minimum lot size of 5000 square feet, p_0 would be higher than the corresponding p_0' under a lower 1400 square feet MLS. However, p_h' would be higher than p_h since any marginal increase in house square footage between 1400 and 5000 square feet would need to be accompanied by more land.

Now let's suppose these costs that the builder faces passes over into the user cost of housing for the household. In a static consumer choice model, this implies that a decrease in the MLS rotates out the part of the budget constraint for the household that is under the old MLS. Hence, assuming that substitution effects dominate, the household's choice goes from X to X', representing a decrease in the household's demand for house size. A colloquial explanation for this phenomenon is: "You might as well build a big house if you have a big plot of land."

In Figure 1.3, I look at households with interior solutions that choose a house size h between 1400 and 5000 square feet. Under a large 5000 square foot minimum lot size (MLS), the slope of the cost graph represents the marginal cost of size, which represents merely the additional materials and construction on the ground floor. However, when the minimum lot size falls to 1400 square

Figure 1.3: Change in MLS leads to Rotating Shift in Budget Constraint

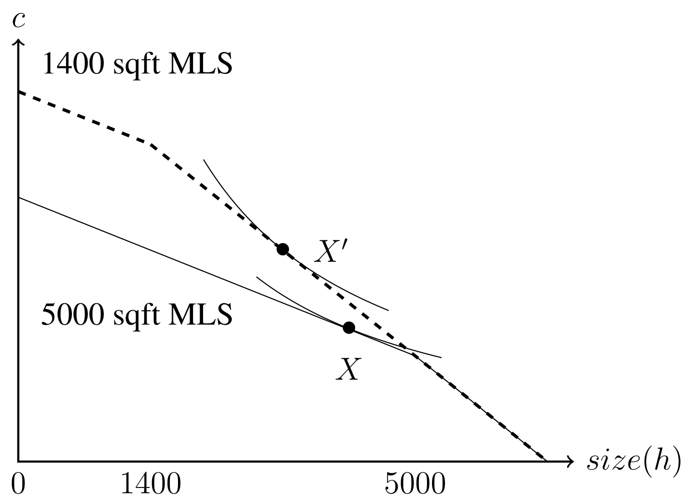

feet, any additional increase in housing size either must be built upwards or must require additional square footage of land. It is this latter channel that increases the slope of the cost curve, i.e., average marginal price of an additional square foot. Hence, in a highly stylized environment where builders can build right to the edge of their lot, a reduction in the minimum lot size from 5000 square foot to 1400 square feet represents a rotation of the budget constraint around a original hypothetical endowment point where the household could have consumed 5000 square feet (and spent the rest on other consumption c).

1.2.1.1 Welfare Heterogeneity

In the simple model illustrated in Figure 1.3, going from a 5000 square foot minimum lot size to 1400 square foot minimum lot size increases welfare for households who consume smaller housing (i.e., below 5000 square feet). The welfare gains comes from two sources:

1. The fall in cost of one's own house, but this gain is smaller the larger the initial demand for house size.

2. The ability to re-optimize and choose a new house size. How these gains change based on income and household size depends on the specific features of preferences.

In essence, these two channels represent a decomposition of the total welfare gains into two parts. The first part is the mere difference between the the dotted budget line and the solid budget line, representing the extra value of consumption if house size could not change; the nature of the MLS makes this amount a perfectly correlated decreasing function of initial house size. However, the ability to re-optimize is important, not only as a significant part of the welfare gains, but because these gains could, in principle, vary across households based on their preferences. In Section 1.4.2, I show the decomposition of the welfare effects and show that the re-optimization gains are economically significant on average but are only weakly positively correlated with income and household size.

1.2.2 Outline of Model Structure

To put the theorized mechanism's intuition into more rigorous terms, and to motivate the difference-in-difference and synthetic control specifications in the next empirical section, I outline an equilibrium model of housing below.

Figure 1.4: Model Structure

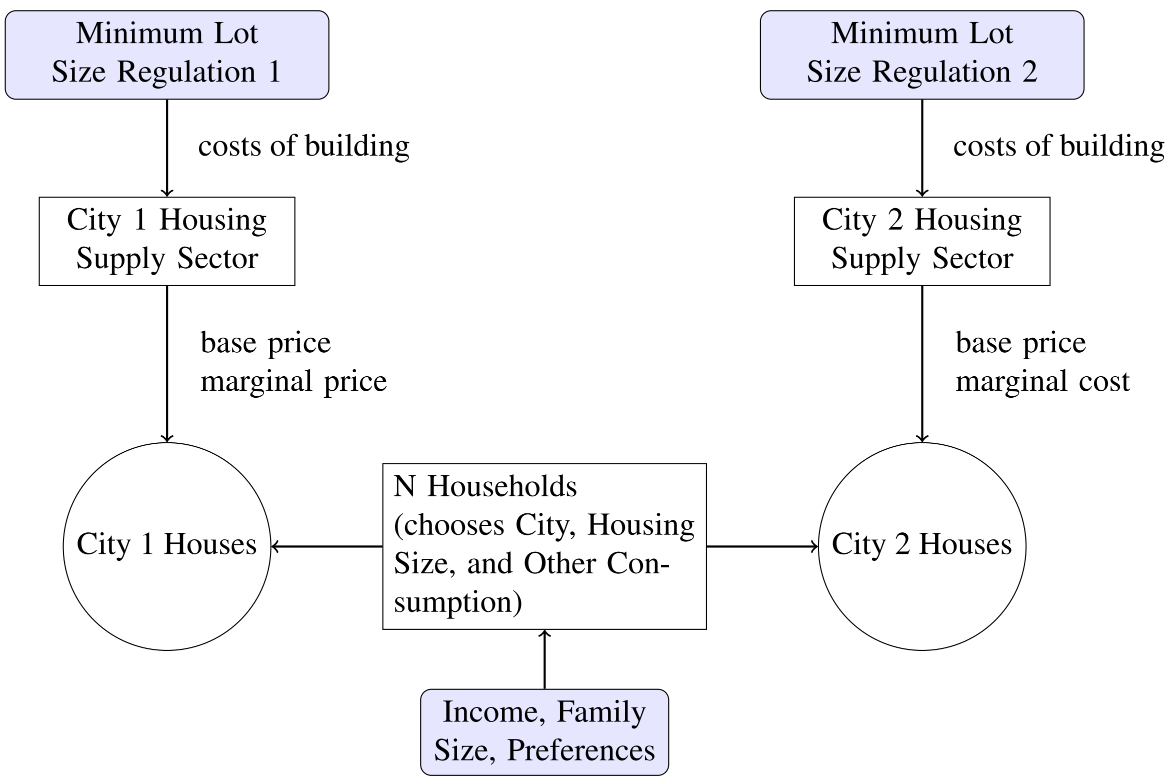

Notes: Sources of exogeneity in model highlighted in light blue.

There are two locations in which households can live. One of these locations will face a change in regulation. Within each location, there are housing supply sectors that face different costs of building different types of housing. The households consist of overlapping generations of people at different points in the lifecycle, and they make decisions about consumption and housing size demand over their lifecycle. The households take prices, income, family size, and idiosyncratic preferences (for housing and for different locations) as exogenous.

1.2.3 Housing Demand

Following the spirit of Aguiar et al. (2021), household i's decisions over their lifetime of length N can be written as:

$$\max_{L} \max_{\{c_{it}\},\{h_{it}\}} \sum_{t=0}^{N} \beta^t \left(\frac{c_{it}^{1-\frac{1}{\eta_c}}}{1-\frac{1}{\eta_c}} + \frac{(\theta_{it}\xi_i h_{it})^{1-\frac{1}{\eta_h}}}{1-\frac{1}{\eta_h}} \right) + \zeta_i^L$$

where

1. c_{it} and h_{it} are the choice variables for consumption and housing size, respectively
2. L represents location
3. θ_{it} is an age/family shifter of size demand
4. ξ_i, ζ_i^L are idiosyncratic preference terms for housing and location
5. η_c and η_h are elasticities

The household is given a lifetime income draw M_i, and they all face the same real interest rate r. The budget constraint in each location L differs. Particularly:

$$\sum_{t=0}^{N} \frac{c_{it} + p^L(h_{it})}{(1+r)^t} = M_i \tag{1.1}$$

where $p^L(h_t) = p_0^L + p_h^L h_t$ is the pricing function for housing in location L. Although it looks like this function has linear form, the estimation allows for nonlinear pricing in the form seen in Figure 1.2. Thus, housing regulations affect each location's household decision through the effects on the pricing function. As explained before in the intuition, a restrictive MLS regulation is expected to increase p_0^L but decrease p_h^L for that particular location (for a large chunk of housing sizes in the middle of the house size distribution).

The structure of these preferences are important in two ways:

$$\frac{d \log h_i}{d \log M_t} = \beta_i = \frac{\eta_h}{\bar{\eta}}$$

where $\bar{\eta} = \sum_i \eta_i \frac{p_i h_i}{M_h}$ First, the difference between η_c and η_h allows for housing to be a necessity good, i.e., expenditure on that good rises less than linearly with income. This is an important stylized fact about housing (and housing size) as described in Appendix A.2.1, and will have important implications for income effects (and therefore welfare calculations at different places in the income distribution).

$$\underbrace{\frac{\Delta \log h}{\eta_h} - \frac{\Delta \log c}{\eta_c}}_{\text{change in relative demands}} = \underbrace{-\Delta \log p}_{\text{change due to prices}} + \underbrace{\frac{\eta_h - 1}{\eta_h} \Delta \log \theta}_{\text{change due to age or family size}} + \underbrace{[\frac{\eta_h - 1}{\eta_h}] \Delta \log \xi}_{\text{change due to idiosyncratic shocks}}$$

Second, as seen in the above equation, these preferences allow for a clear channel in which demographics and prices both shift the demand for house size. The change due to prices represents a substitution effect. Any welfare gains or losses due to change in prices will interact with both the budget constraint, as well as the demographic demand shifter. In the end, there will be a quantitative analysis of the total welfare effect across different demographics.

Family size and age enter into the preference term θ_{it} structurally.[3] Specifically:

$$\theta_{AZ} = \alpha_0 + \alpha_1 A + \alpha_2 A^2 + \alpha_3 Z \tag{1.2}$$

where α_n are parameters of the age curve, A is the age of the household and Z is the household (family) size. There are priors about the sign of these parameters based on existing empirical and theoretical work. First, stylized facts, like in Banks et al. (2015), about the hump-shaped demand for housing over the lifecycle suggests α_1 is positive and α_2 is negative. Second, the strong positive correlation between house size and household size suggests $\alpha_3 > 0$. As will be seen in the estimation section, these priors are confirmed when the model is estimated.

Permanent income M_i is measured as a weighting w between current income Y_i and average income within the education/industry group $\bar{Y_g}$

$$M_i = G[wY_i + (1 - w)\bar{Y_g}] \tag{1.3}$$

where G is a multiplier to convert annual income to lifetime income. In the model estimation, G is set so that lifetime income is simply the lifetime value of the implied annual income, given assumptions about an interest rate and some growth rate of annual income. The weighting between household idiosyncratic income and the education/industry group's income is important for the estimation: by averaging each particular household income with its group average, model fit is substantially improved, likely because housing decisions are based on permanent income and therefore year-to-year household income is less informative than education and industry.

Given a joint distribution of income, preferences, age, and family size, the total demand for

[3]Several different functional form of this equation were used. The form used in this baseline represents a tradeoff between having a relatively few parameters to estimate and having relatively good fit. Additional interaction terms did not significantly improve model fit.

housing at any given time is simply the resulting full distribution of housing sizes that are the solutions to the household problem.

1.2.4 Housing Supply

Given that housing demand is a distribution of sizes, housing supply also consists of a distribution of house sizes. As a simplifying assumption, I break apart the supply distribution into a discrete number of different housing sizes q.

A competitive housing supply sector must be indifferent between producing each type of housing; otherwise, more of that housing will be built. This can be motivated by thinking of a representative firm that chooses housing investment each period to maximize joint profits

$$PROFIT(\{I_{1t}\},\{I_{2t}\},...\{I_{Qt}\}) = \sum_{k=0}^{\infty} \frac{\sum_q p_{q,t+k}H_{q,t+k} - P(I_{t+k})}{(1+r)^k}$$

where

$$\begin{aligned} I_t &= \pi_1 I_{1t} + \pi_2 I_{2t} + ... + \pi_Q I_{Qt} \\ P(I_t) &= \sigma(I_t)^\gamma \\ H_{qt} &= (1-\delta)H_{q,t-1} + I_{qt} \end{aligned}$$

Here, each q type of housing is a stock that depreciates at rate δ, but is replenished by new investment I_t^q. The quantity flow of housing that the stock produces rents at rate $p_t(q)$. The cost of total investment I_t is given by a convex function $P()$, which has a form with a steepness parameter σ and convexity parameter $\gamma > 1$.

For all q, the first order conditions for profit maximization give the following optimal investment decision each period:

$$\underbrace{\log(\sigma\gamma) + (\gamma-1)\log I_t}_{\text{common across types}} = \underbrace{\log PV_{qt} - \log \pi_{qt}}_{\text{type specific}} \tag{1.4}$$

where $PV_{qt} = \sum_{k=0}^{\infty}(\frac{1-\delta}{1+r})^k p_{q,t+k}$ is the present value of the stream of future rents for housing type q. In a steady state, given $p_{q,t+k} = p_q$ the asset price of a house at any given time is proportional to its rent:

$$PV_q = \frac{1+r}{r+\delta} p_q \tag{1.5}$$

Hence, by differencing two levels of q, $\Delta \log PV = \log PV_q - \log PV_{q'} = \Delta \log p = \Delta \log \pi$.

That is, anything that identifies changes in log asset prices also identifies changes in log rents, which also identifies changes in the marginal costs of building. Hence, overall supply and demand may result in shocks to overall housing prices across different areas, but within an area, the relative costs (in logs) of different housing types remain completely determined by their relative costs of building; hence, in a difference-in-difference setting, the marginal costs of building "pass-through" to equilibrium prices.

1.2.5 Equilibrium

Given a distribution of preferences (over housing and location), income, family size, and costs of building, an equilibrium is a set of prices in which, in each location L, (a) the demand distribution is consistent with the household problems (b) the supply distribution is consistent with the builder's problem (c) the supply distribution equals the demand distribution.

1.2.6 Model Summary

The basic model outline is a model of two locations where each housing supply sector's costs of building are influenced by regulations. These costs are passed through to households in terms of their user cost of housing. A mass of households makes decisions about where to live and the types of housing size needed in the city that they live in. At the core of the model is the non-homothetic preference structure, which captures a key feature of the necessity good aspect of housing consumption and allows for a range of income effects across the income spectrum. The same elasticity η_c and η_h parameters that govern the nature of the income effects also govern the overall housing share and the magnitude of the substitution effect in response to changes in the marginal price of house size.

The only question remaining is, how exactly do the housing supply sectors' costs shift? Given our theoretical framework for the minimum lot size regulation, I turn to the causal and reduced form evidence from the deregulation event in Houston to estimate both the compositional effects and price effects of changing a minimum lot size. Afterwards, the estimated magnitude of those effects are fed into the model as a policy experiment, and the distribution effects of the experiment will be reported.

1.3 Empirical Analysis - Houston

To understand the effects of a minimum lot size regulation, I study a minimum lot size change in Houston's central area that was enacted in 1998 and implemented in 1999. The seminal reference for this policy change is Gray and Millsap (2020). The 1998 reform did one main thing: it reduced the minimum lot size in most of Houston's inner-ring (within I-610) area to as low as 1400 square feet. This area represents the central business district of Houston as well as its surrounding neighborhood, but does not include the outer areas that are close to Houston's borders with its suburbs.

The actual minimum size in each Houston block could have depended on a variety of factors, from open space requirements and community opt-out at the neighborhood level. Because of heterogeneity across neighborhoods in terms of the intensity of treatment, I view the estimated effect as an intent to treat.[4] Gray and Millsap (2020) has shown that this deregulation event spurred development of many smaller lots in middle income neighborhoods. Figure 1.5 shows a count of smaller lots (< 5000 square feet) developed.

Figure 1.5: Number of Smaller Lots (< 5000 SQFT) Developed in Houston

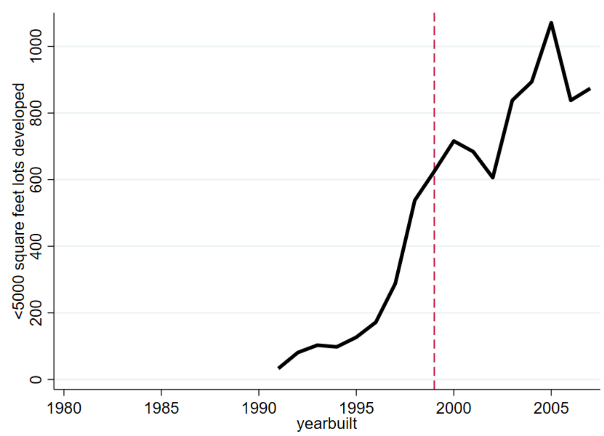

The larger context is that Houston has always relied on a variety of other regulations (like private covenants and regulations on parking) to plan and control development. One of the main regulations was the minimum lot size. Residential lots were, at least on the books, required to be at least 5000 square feet. Deviations from this regulation were relatively rare because they required variances (special approval from the planning department).

The reform happened in large part due to a community desire for urban renewal and the willingness to attract young professionals into the area. However, there was substantial opposition, largely from existing homeowners, who had a desire to maintain the characteristics of their neighborhood and stabilize housing prices. Gray and Millsap (2020) argue that the reform was possible because of a grand political compromise: this political innovation allowed individual blocks or small neighborhoods to opt-out of the reform, which assuaged much of the opposition. Overall, this policy

[4] Attempts to exclude the census tracts where there were the largest numbers of opt-outs to the policy did not change the results qualitatively.

reform was seen as a pioneering change. A crucial assumption in identification is that the political reform was not unique to Houston in such a way as to break the parallel trends assumption. Extra care, therefore, is taken to create a synthetic control which looks very much like Houston in the pre-period before the reform.

Note that a 5000 square foot lot was significantly larger than the median house in Houston, and the extra land required for a housing unit likely increased housing costs. The images in Figure 1.6 show an example of lots that were developed before the 1998 reform: many single family houses spanned a part of a large plot of land. A significant amount of the lot was used as backyard or landscaping.

Figure 1.6: Example of Pre-reform Housing in Houston: Built in 1995, 2500 square feet living space, 10000 square foot lot

Note: Images from Google Maps and Google Street View, edited to anonymize street number.

I also show an example of post-reform housing in Figure 1.7. This block was likely subdivided into four smaller 2600 square foot lots. There are several notable characteristics; namely, the square footage of these houses are much smaller. Moreover, they were built with much less green space and were much closer to the limits of their lots.

Figure 1.7: Example of Post-reform Housing in Houston: Built in 2004, 1500 square feet living space, 2600 square foot lot

Note: Images from Google Maps and Google Street View, edited to anonymize street number.

The question that I seek to empirically answer is not confined to the overall effect on the quantity of lots developed. Rather, I use the model to investigate the mechanisms of how welfare is affected, which types of households are disproportionately affected, and the types of selection that would be predicted from such a model. To connect the model to the data, I estimate the effect of the deregulation event on the average housing built each year, as well as on the marginal price of an additional square foot. These reduced form causal estimates will then be fed back into the model to evaluate the mechanisms and welfare effects on different demographics.

In the following sections, I first describe the data used to conduct the analysis of Houston. Secondly, I describe the empirical models, including both traditional difference-in-difference estimators and synthetic control methods to describe the effects on Houston.

1.3.1 Data

The main dataset used in this paper is the deed and tax data on housing characteristics as collected by Corelogic. These are property tax records from different jurisdictions that have been compiled into one proprietary dataset. The main benefit of this dataset is that precise coordinates of the house are available. Also, for the majority of jurisdictions, there are precise measures of square feet of floor space. Accompanying this dataset is the transactions data compiled from deed records: this additional dataset contains sales transaction data, in terms of date and price, for houses that are linked to the housing characteristics data.

The model section makes use of the public use version of the long form Census (from 2000) provided by IPUMS.[5] Parts of the Appendix use this same data from 1960 to 2010 to document historical trends. The long form Census is a representative and comprehensive subsample of American households and has a section on dwelling characteristics. The main variable of interest is the reported number of bedrooms in their house. Other variables of interest are household income, education, age, and household size. With this rather complete dataset, I am able to estimate and calibrate a model that features a joint relationship between income, age, and fertility.

I also use the American Housing Survey, which is a survey of housing starting in 1975 but has no representative size measures of housing in square feet until 1985. This survey data is used in this paper to check the larger Census data, to verify house size trends in the Corelogic database, and as a way to get a back-of-the-envelope calculation on the changes in square feet for each additional room or bedroom.

1.3.2 Empirical Model: House Size

1.3.3 Synthetic Control

Given the presence of one treated unit, the need to satisfy parallel trends assumptions in the comparison group, and given the large pool of possible comparison cities to Houston, I use synthetic control methods to estimate the effect of the policy change on Houston's size of new housing built. The intuition behind the synthetic control methods is to combine a matching estimator with the

[5]Data from Steven Ruggles, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. IPUMS USA: Version 9.0. Minneapolis, MN: IPUMS, 2019. https://doi.org/10.18128/D010.V9.0

difference-in-difference framework, so as to narrow the set of comparison cities and "synthetically" create a hypothetical comparison city that would satisfy the parallel trends assumption.

The synthetic control method follows Abadie et al. (2010) and subsequent work by Abadie (2021). In the spirit of their work, I choose a collection of donor pool cities that are plausible (and potential) comparison cities to Houston, i.e., "donor cities". Having reasonable choices for donor cities is important in avoiding problems with extreme interpolation and overfitting. As such, I restrict the set of cities to be within the areas in the United States around Texas, with the addition of cities in Florida and Georgia. Overall, I have a donor pool of 29 cities.

For donor cities, I use all available jurisdictions in the Corelogic data that satisfy the following conditions:

1. Jurisdiction is in Texas or nearby states in the South (Texas, Oklahoma, Arkansas, Louisiana, New Mexico, Florida, Georgia).

2. Jurisdiction has at least 350 units built per year from 1991 to 2007

3. Jurisdiction has sales data throughout the above time period

This filtering selects on comparable geographies in the region, as well as both data availability and the size of the jurisdiction. The characteristics of the donor cities in comparison to Houston is given in Table 1.1.

Table 1.1: Summary Statistics for Houston vs Donor Cities

Variable	Houston	Donor Cities
Minority Population (mean)	63.9%	44.3%
Median HH Income (mean)	73272	78470
Poverty Rate (mean)	11.3%	8.5%
Median Rent (mean)	1088	1124
MSA Pop Growth (1991-1997)	15.2%	19.4%
Density (1997)	3371.7	2579
Jurisdictions	1	29

Note: Means for first four variables are observed at the census tract level. These values are equivalent to weighted averages of census tract characteristics, where weights are determined by observations of residential housing units in the Corelogic dataset.

1.3.3.1 Methodology: Theory, Predictor Variables, and City Weights

Synthetic control methods vary in style. In the end, the methods all choose a convex combination of comparison cities from the "donor" pool to create a synthetic comparison city. The weights for each city are chosen to minimize a given norm of distance of a vector of predictor variables between Houston and other cities. Consistent with common practice, I use the weighted mean square prediction error (MSPE) of the outcome variable in the pre-period as a norm.

Let each j "predictor" variable X_{ij} be an observable associated with a city i. Let $X_j^{houston}$ be the j variable associated with the treated city (Houston). The MSPE is given by:

$$\frac{1}{J}\sum_{j=1}^{J} v_j\left(X_j^{houston} - \sum_i w_i X_{ij}\right)^2$$

where J is the number of predictor variables, w_i is the associated weight for each city and $v_j > 0$ is a separately estimated (or exogenously given) weight for each predictor variable. Note the normalization restrictions: $\sum_i w_i = 1$ and $w_i > 0$, the latter which eliminates synthetic controls which arise from extrapolation. Also, note that the j variables can include both time-varying and time-invariant predictor variables.

The weights on each variable v_j are important because they also determine the optimal weights w_i for each city chosen to be in the synthetic control. To reduce idiosyncratic biases introduced by the researcher in their own personal choice of variable weights, I use the standard choice of the variable weight vector V. Specifically, the weights V are chosen to minimize the following MSPE of the outcome variables in the pre-period (1991-1998).

$$\frac{1}{8}\sum_{t=1991}^{1998}\left(Y_t^{houston} - \sum_i w_i(V) Y_{it}\right)^2$$

where Y are the outcome variables (log of average square feet of new housing built) and $w_i(V)$ are the estimated optimal weights conditional on a choice of V. As such, the two previous equations, used as objective functions, define a nested minimization problem for both variable weights and city weights. Estimated variable (predictor) weights are reported in Appendix A.4.1.

The predictor variables (the j variables) used follow in the spirit of Abadie et al. (2010) in the use of a combination of evenly spaced (i.e., skipping years) outcome variables and other predictor variables. This is a compromise between competing styles in the literature. For example, Ferman et al. (2020) emphasize the importance of matching on a large number of pre-period treatment outcomes to avoid specification searching, while Cavallo et al. (2013) practice limiting the matching to a few pre-period outcomes as a test of out-of-sample validity. Specifically, I use the odd-yeared lagged outcome variable (log average square feet) during the pre-period, augmented with variables that describe the population, income, and price characteristics of Houston. Overall, the synthetic control city matches Houston very well in terms of MSA-level population growth, median rent, and odd-year outcome variables. However, Houston remains more minority, poorer, and more dense than the synthetic control city. In Appendix A.4.2, I explore alternative specifications where matching is done on the outcome variable for every pre-period year, and where different variables are dropped. In general, I show that the results are robust to choices of predictor variables.

Table 1.2: Predictor Variables for Synthetic Control

Variable	Houston	Synthetic Control City
Minority Population	63.9%	45.3%
Median HH Income	73273	89270
Median Rent	1088	1165
Log Square Feet (1991)	7.838	7.830
Log Square Feet (1993)	7.801	7.806
Log Square Feet (1995)	7.810	7.818
Log Square Feet (1997)	7.831	7.826
MSA Pop Growth (1991-1997)	15.2%	15.3%
Density (1997)	3371.7	2860.8

Note: Minority Population, Median Income, and Median Rent characteristics are tract-level characteristics weighted by housing units built in the pre-period.

The weights and cities in the synthetic control are given in Table 1.3. There are two notable observations: first, virtually all weight is on cities in Texas, suggesting that the matching algorithm may be picking up regional-year fixed effects unique to Texas (i.e., not present in other major cities in the donor pool like Atlanta or Orlando). Second, they are some of the larger cities in Texas, which are natural and intuitive comparison cities to Houston.

Table 1.3: Baseline Model: Estimated Synthetic Control Weights

City	Weight
Plano, TX	0.388
San Antonio, TX	0.309
Austin, TX	0.237
Sugar Land, TX	0.064
Tulsa, OK	0.002

1.3.3.2 House Size Results

In Figures 1.8 and 1.9, the bold red line indicates the log average square feet of new housing in Houston relative to the synthetic comparison city. Qualitatively, the Houston trajectory of log square feet matches its synthetic control prior to the deregulation event. After 1999, however, the two series start to diverge. The relative change of Houston's from the pre-period average to the average around 2005-2007 is a decrease of house size of about 12.5 log points.

Figure 1.8: Houston vs Synthetic Control Outcome, Minimum Lot Size Reduction in 1999

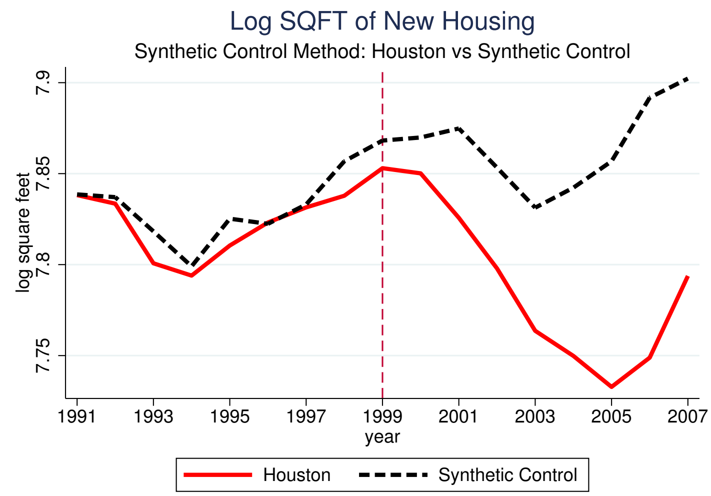

Notes: The synthetic control method chooses a convex combination of control cities that minimizes an distance function of variables from the cities.

Figure 1.9: Houston: Synthetic Control, Minimum Lot Size Reduction in 1999

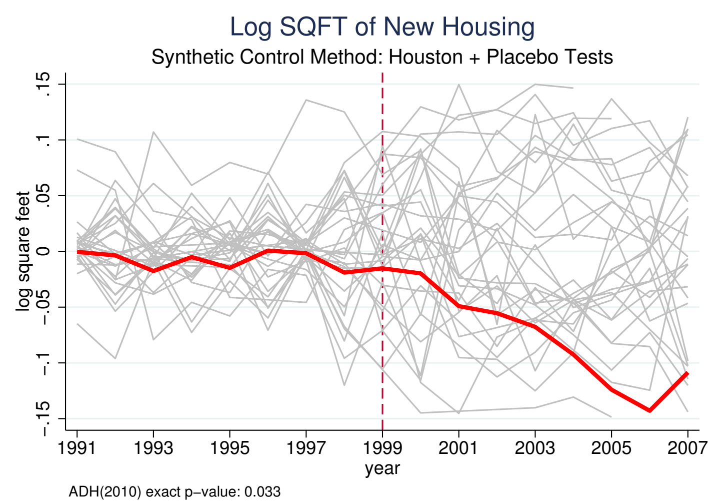

Notes: The synthetic control method chooses a convex combination of control cities that minimizes an distance function of variables from the cities. Gray lines are results of placebo tests where the same synthetic control procedure is repeated for all cities in the donor pool.

Figure 1.9 plots the differences between Houston's trajectory and that of its synthetic control in bold red. The plot also includes gray lines which are placebo tests: they represent the trajectory of every other city in the donor pool, relative to a synthetic control which is generated by the same matching algorithm. One important consideration is that much of the noise results from donor cities which have poor matches in the data (i.e., perhaps they are extreme in their predictor variables and are hard to interpolate with a convex combination). However, one can still visually see that by 2005-2007, Houston is a relative outlier in terms of the magnitude of the decline in average house size of new housing built. However, visual interpretation of the significance of Houston's results may be unreliable because identified large effects in the post-period for donor cities may be caused by poor fit in the pre-periods for those cities; consequently, a more rigorous way of doing statistical inference is used.

The standard way to do rigorous inference about the significance of the effect sizes found in Houston is to calculate the ratios of Root Mean Squared Prediction Errors associated with post-period to pre-period effects. Specifically, the test statistic for city i is:

$$RATIO(i) = \frac{\sqrt{\sum_{t \in post} \left(Y_t^{houston} - \sum_i w_i Y_{it} \right)^2}}{\sqrt{\sum_{t \in pre} \left(Y_t^{houston} - \sum_i w_i Y_{it} \right)^2}}$$

Intuitively, outcomes of cities that diverge significantly from their synthetic controls, relative to that divergence before the treatment, show more statistically relevant results. The exact Fischer p-value is therefore the rank of this ratio (amongst all donor cities) as a fraction of the total number of cities. Because Houston's calculated ratio is the highest out of 30 cities (Houston in addition to 29 donor cities), 0.033 is the calculated p-value. More noteworthy, it is possible to examine the distribution of these ratios to see whether Houston's test statistic ratio stands out.

Figure 1.10: Statistical Significance of Houston's Synthetic Control Results

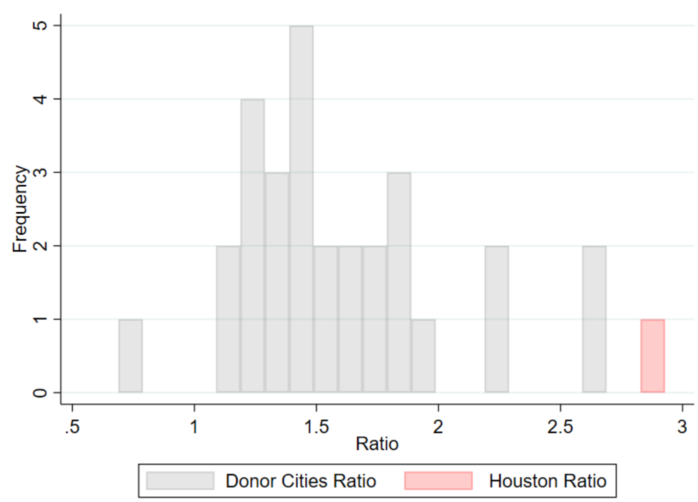

Notes: This is a histogram of the calculated test statistic ratios of RSMPE (post) to RSMPE (pre) amongst all donor cities and Houston. Houston's ratio is an outlier.

To check the robustness of these baseline results, I turn to two methods: first, I run an alternative synthetic control model with only census-tract level predictor variables (related to city characteristics like poverty rate, rents and income) and odd-year outcome values. I also run a version which only matches on outcome variables in the full pre-period sample. I show that the results do not substantially change (see Appendix A.4.2). Finally, I use the traditional difference-in-difference estimator for all cities in the donor pool and show that the resulting magnitude of the size estimates, even though they exhibit some pre-trends, are squarely consistent with the baseline of about a 10-15% reduction in new housing size (see Appendix A.4.3).

1.3.3.3 Difference-in-difference Entire Distribution

Figure 1.11: Synthetic Control Cities vs. Houston

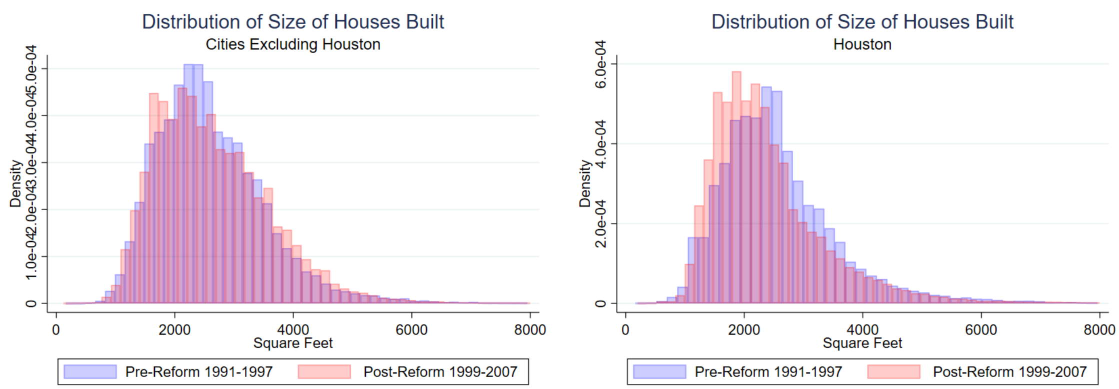

Here, I show the effects on the entire distribution of new housing built before and after the MLS deregulation. In the (unweighted) synthetic control cities outside of Houston, the distribution of pre-reform and post-reform housing size looked approximately the same, with possibly some mass moving to the right of the distribution. However, there is a clear visible shift of mass when the same graph is shown for Houston. This shift in distribution appears to happen throughout the areas where the distribution has substantial support. There is no bunching visible in any area. In other words, the direction and nature of this shift is consistent with the theorized mechanism of a change in the price of house size.

1.3.3.4 Floor Area Ratios

I show more evidence that the Houston 1998 reduction in the minimum lot size is consistent with the idea that the extra lot size posed extra costs for many households. To illustrate this, I run a difference-in-difference regression weighted using synthetic control city weights, looking at Houston's floor space area ratio (FAR) relative to the synthetic control cities. I use several measures of the FAR: the level itself, the log level, and an indicator for when the FAR exceed 50% of the lot size. I find that after the policy change, Houston's new houses used a statistically significant larger percentage of their lot space than before, relative to its synthetic control city. On average, FAR

Table 1.4: Floor Area Ratios

Dependent Variable: Different Measures of Floor Area Ratio)	(1) FAR	(2) Log FAR	(3) $\mathbb{1}(FAR \geq 0.5)$
Houston*Post	0.060	0.116	0.048
	(0.009)	(0.053)	(0.011)
Observations	280433	280433	282876
City FE	x	x	x
Year FE	x	x	x
Synthetic Control Weights	x	x	x

Standard errors in parentheses

increased about 6 percentage points. That presents about a 11 log point increase. The third regression specification suggests that a significant proportion of this rise was the shift to FARs higher than 0.5.

1.3.4 Empirical Model: Price of House Size

1.3.4.1 Basic Price Regression

To analyze the direct effect of regulations on Houston's marginal cost of house size, I use the following empirical model that identifies the differential effect of the 1998 policy change on prices. Because there are census tract fixed effects, the identifying assumption is that the square footage of any given house (and its interaction with the Houston jurisdiction), conditional on being in the same census tract, is uncorrelated with unobservables that affect house prices. To be clear, the mere fact that unobservables (like granite countertops) are likely correlated with house size is not necessarily a problem, as long as these correlations are stable across space and time. As such, the difference-in-difference nature of this empirical model is capable of differencing out such possible biases.

Price Regression: p_{ijct} is the sales price for house i, census tract c, jurisdiction j, and sales year t.

$$p_{icjt} = \zeta_c + \eta_t + \sum_t \lambda_t(SQFT_{icjt}) + \sum_t \alpha_t * houston +$$
$$+ \sum_t \beta_t(SQFT_{icjt}) * houston + \epsilon_{icjt}$$

The parameters of interest are α_t which represents Houston's relative "base" price, and β_t which represents Houston's relative price per square feet. The theory on the effect of the minimum lot size predicts that α_t will decrease due to the deregulation and β_t will increase.

Because this full specification requires the estimation by year by year coefficients which are noisy, I run a simplified version of the regression above by pooling all the years into either pre-period or post-period indicators. Looking at the standard errors on the coefficients is essentially running a statistical test of whether the individual β_t and α_t coefficients, averaged over pre and post periods, are statistically different. Slightly abusing notation, in Table 1.5, I denote the pooled coefficients for the post period relative to the pre-period as β and α. I run specifications with and without individual year fixed effects (instead pooling the pre/post periods' effects), crossed with different comparison groups (all donor cities vs. restricted to the synthetic control cities as identified in the previous section).

1.3.4.2 Price Results

Table 1.5: Price Regressions for New Housing

Dependent Variable: Sales Price (2010 dollars)

	(1) Price	(2) Price	(3) Price
Houston*Post (α) (in thousands)	-27.721	-12.404	-11.924
	(7.022)	(3.866)	(3.079)
Houston*Post*SQFT (β)	24.860	20.917	22.266
	(6.195)	(3.827)	(4.356)
Observations	199996	71776	71776
Tract FE	x	x	x
Year FE	x	x	x
Synthetic Control Weights			x
Comparison Group Sample	Donor Cities	Synthetic Control Cities	Synthetic Control Cities

Standard errors in parentheses

The regression results in Table 1.5 show that Houston had a significant relative increase in price per square feet between the two periods (i.e., this price increased after the decrease in the minimum lot size). The relative increase is substantial; about 22 dollars per square foot, an estimate which does not fluctuate too much across the different specifications. In log terms relative to the marginal price per square foot in the Houston pre-period, this is an **increase of about 14 log points** (i.e., $\Delta \log p \approx 0.14$). The preferred specification (third column) suggests that the overall base price of housing decreased about \$12000; these intercept estimates, however, are relatively noisy across different specifications compared to the marginal price (slope) estimates, but they are all consistent with a decrease in the base cost of a house.

A more relevant test of the price mechanism is to directly show whether the lot size channel is responsible for a large proportion of the marginal cost of housing floor space. I do this by controlling for the size of lot on which each house sits.

$$p_{icjt} = \zeta_c + \eta_t + \sum_t \lambda_t(SQFT_{icjt}) + \sum_t \alpha_t * houston +$$
$$+ \sum_t \beta_t(SQFT_{icjt}) * houston + \underbrace{\sum_t \chi_t LotSQFT_{icjt}}_{\text{lot size controls}} + \epsilon_{icjt}$$

Controlling for the size of the lot should control for any marginal cost changes that are assigned to marginal floor space changes. Indeed, the results in Table 1.6 are consistent with that. After the inclusion of lot size controls, the estimated change in the marginal price of house size is significantly lower. In the preferred baseline specification (weighted Synthetic Control Cities sample), the relative change in that price for Houston is not even statistically distinguishable from zero.

Table 1.6: Price Regressions for New Housing, Controlling for Lot Size

Dependent Variable: Sales Price (2010 dollars)

	(1) Price	(2) Price	(3) Price
Houston*Post (α) (in thousands)	-28.122	-13.799	-9.691
	(6.423)	(5.068)	(5.821)
Houston*Post*SQFT (β)	7.231	4.559	2.775
	(2.237)	(1.789)	(1.509)
Observations	199371	71762	71762
Lot Size Controls	x	x	x
Tract FE	x	x	x
Year FE	x	x	x
Synthetic Control Weights			x
Comparison Group Sample	Donor Cities	Synthetic Control Cities	Synthetic Control Cities

Standard errors in parentheses

There are several main takeaways from the empirical results from Houston. The first is that the average size of new housing decreased by about 12.5 log points after a policy change which decreased the minimum lot size from 5000 square feet to 1400 square feet in most of the areas of Houston. The second is that the price effects are consistent with the hypothesized mechanism: the marginal cost of an additional unit of floor space increased (about 14 log points) because of the additional lot size needed to complement the house; after controlling for lot size, the change in this cost is not statistically different (relative to other cities). Finally, the ratio of identified house size and price effects imply an elasticity of $\epsilon = \frac{-\partial \log(h)}{\partial \log p} = \frac{12.5}{14} = 0.893$ which is used to verify the fit of the model in the next section.

1.4 Quantitative Analysis

The objective of doing a simulation analysis of a Houston-like deregulation is to (a) understand the welfare implications of the minimum lot size policy experiment (b) generate other testable predictions about location selection that are related to the main channels being analyzed. In the following section, I present details about how parameters are being calibrated and estimated. I present some results about model fit. Then I detail the exact experiment being run and show the results, as well as the intuition behind such results. Finally, I discuss the selection mechanisms and see whether or not they are confirmed in the data.

1.4.1 Model Estimation and Fit

The model parameters are either chosen to match plausible values, or they are estimated to match the cross-sectional patterns of house size choice in the public Census. To accomodate the overlapping intergeneration structure, I allow each household to live for 6 periods (each period representing 10 years) starting from age 25. The simulation of the model has two thousand households per generation total with a distribution sampled from the 2000 Census. Each generation is then weighted to match the population statistics of the 2000 Census cross-section.

The estimation procedure is to match both (1) average house size demand by age and quintiles of hh-size/income and (2) log variance of demand. These model averages are estimated by a two-step feasible SMM estimator, which minimizes a weighted quadratic of the difference between model generated moments and data moments. I list the averages/variances targeted in Table 1.7.

Table 1.7: Moments

Moments	Description	Associated Parameters
Average house size demand	150 averages by age, quintiles of family size x income	α, η
Log variance of house size demand	1 population statistic	κ

In Table 1.8 are the main parameters of the model. The interest rate r is set at 10% per decade. The permanent income multiplier G is derived from a growth rate of 3%. The income weight w is set so that 90% of permanent income is based on the household head's education and industry, and only 10% is based on current income. This captures the effects of a conditional mean reverting process income where idiosyncratic income converges to the group average over time.

The housing size parameter, demographic shifters, and variances are estimated from the cross-sectional variances. What is notable is that the estimated parameters, disciplined by the data, speak clearly about the shape of the demographic curve over the lifecycle: housing size needs are highest in the middle and end of the lifecycle, and smallest when households are youngest. Finally, the housing size parameter is significantly smaller than the consumption parameter, which means housing size demand is decreasing as a percentage of income, as income increases. Hence, the estimation disciplines housing to be a strong necessity good. All of these features capture important variance/covariance relationships in the data.

Table 1.8: Parameters of the Model

Parameter	Description	Method	Value (Std Error)	Discipline
η_c	Consumption parameter	Estimated	2.46 (0.058)	Consumption Share
η_n	Housing Size parameter	Estimated	1.38 (0.012)	Engel Curve
κ	Variance of $\log \xi$	Estimated	1.24 (0.070)	Dispersion of Demand
α_1, α_2	Age Shifters	Estimated	85.4 (5.3), -4.78 (0.31)	Age Curve
α_3	HH Size Shifters	Estimated	142.6 (9.6)	Family Size House Demand
r	Interest Rate	Calibrated	10% per decade	
G	Permanent Income Multiplier	Calibrated	106.34	3% growth per year
w	Income Weight	Calibrated	0.1	

Note: Estimated parameters are estimated with the method of simulated moments, using a weighting matrix derived from estimates of the inverse variance of the moments via bootstrapping. Standard errors (in parenthesis) are calculated numerically.

The intuition behind the discipline for the η parameters is that they are estimated to jointly explain the average housing or consumption share in the data. To separately identify η_c from η_h, the estimation procedure is implicitly targeting the curvature of the Engel Curve (in addition to the overall housing share), which is consistent with what is noted in Equation 1.2. The demographic shifters are in vector α that govern the relationship between age and house size demand, as well as family size and house size demand. The second α parameter being positive means that housing demand is increasing overall in age, but the third α parameter denotes concavity of that function, which is consistent with both theory and empirical observation (that housing demand is increasing in young age and then flattens out) . Finally, the last positive parameter in α denotes that housing demand is increasing in family size. Hence, the estimated parameters, as disciplined by the data, generate a model that has structural relationships in directions that are consistent with my priors as informed by the literature and by economic intuition.

1.4.1.1 Model Fit

1.4.1.2 Targeted Model Fit

Next I discuss the quantitative fit. In Figure 1.12, I plot the model generated moments against the empirical moments estimated from the Census data, excluding the moment associated with

Figure 1.12: **Baseline Model:** Simulated Model Averages vs. Empirical Averages for Each Age / HH Size / Income Group

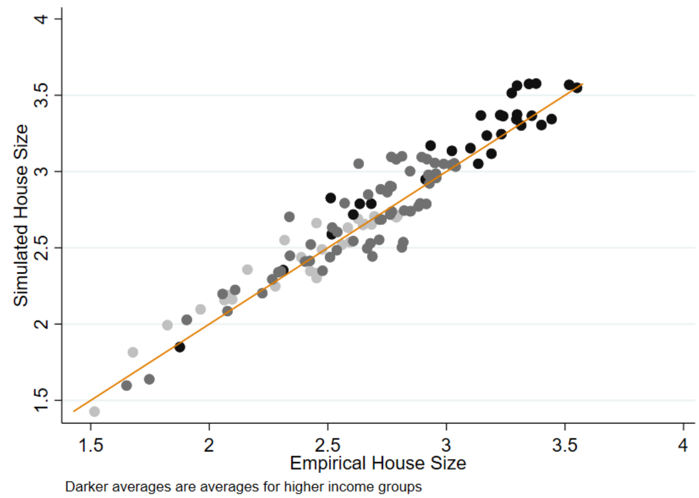

variance.[6] These are essentially conditional averages at different bins of the age, income and family size. The line shown is the 45 degree line, which represents a perfect fit of data and model. The model qualitatively does very well with relatively few estimated parameters.

Figure 1.13: **Alternative Homothetic Model:** Simulated Model Averages vs. Empirical Averages for Each Age / HH Size / Income Group

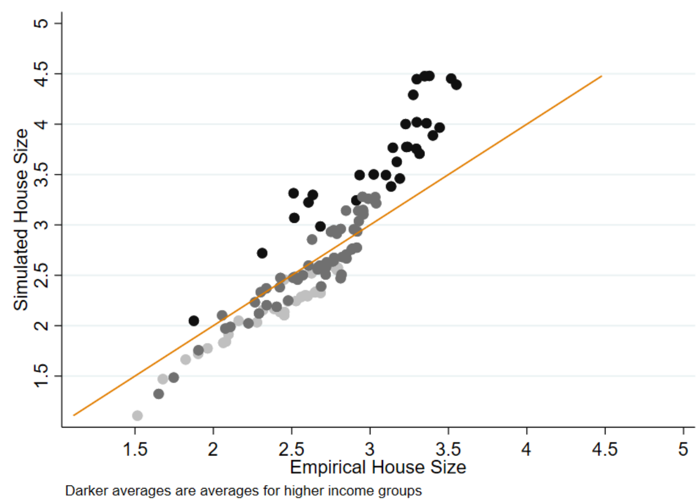

I further explain my contribution of using these preferences to capture the nonlinear Engel curves of housing. By breaking up the model fit line into lower (below 20th percentile) vs middle

[6]The empirical variance and model variance are very close, but their values are on a different scale so they cannot be represented well in the graph.

(20-80th percentile) vs. higher (above 80th percentile) income groups, one can see that model fit is qualitatively the same for different income groups. To illustrate the alternative of using standard homothetic preferences, I re-estimate a restricted version of the model where the restriction $\eta_h = \eta_c$ is imposed; this effectively makes the Engel curves linear and the preferences homothetic. It is natural to expect a loss of model fit with a decrease in the degrees of freedom. What is noticeable in the model fit shown in Figure 1.12, however, is the systematic way in which the model overestimates housing size demand for higher income groups and understimates them for lower income groups. I include this to highlight the dangers of using homothetic preferences and the contribution of using preferences in the baseline model that more accurately capture the necessity good features of housing size demand.

1.4.1.3 Untargeted Model Fit

A stronger test of a model is whether it can explain patterns and features of reality that are not imposed by the researcher. The main object of interest in this model is the elasticity of house size with respect to its price, i.e., $\epsilon = \frac{-\partial \log h}{\partial \log p}$. **This simulated elasticity is calculated to be 0.926**, a number derived by calculating how large a price effect is needed to rationalize the size effect found in Houston. This implied elasticity in the model, however, heuristically only comes from estimates of η_c and η_h based on data related to the housing share and the shape of the Engel Curve; it takes no information about how sensitive demand of housing size is to changes in prices.

The estimated elasticity of 0.926 is very close to the elasticity of 0.893 estimated from the Houston natural experiment in Section 1.3.4.2. The latter is arguably an exogenous policy change that induced an exogenous price change. Hence, it was a more direct way of estimating the relevant elasticity. I argue that the fact that these two drastically different methods agree is a solid confirmation of one aspect of the model to reflect reality.

1.4.2 Simulation Experiment

The experiment run in the model section is to start the model in steady state using the parameters chosen or estimated. In this steady state, prices are constant, and the population distribution is unchanging over time. Consequently, each household lives in their preferred city, optimally choosing housing size and consumption.

The simulated experimental shock is as follows: An increase in the marginal price of size of building π_q is introduced in one city. For such a price change to reflect a realistic minimum lot size deregulation, it has to be neutral near the previous minimum lot size. That is, given that a 5000 square foot minimum existed, the cost of building a house at around 4500 square feet (leaving some area for green space and other purposes) should be about the same before and after the deregulation. Since the Census data is calibrated on a bedroom measure of house size, the neutral pivot point is calculated from an empirical average relationship between square footage and bedrooms. This pivot point is 4.5 bedrooms. Note that this is simply a translation of units, which will be converted to log point changes in the experiment. It in no way suggests that there will not be differing intensive margins of bedroom size available.

1.4.2.1 Long Term vs Short Term

One of the key inputs into the model is the change in prices. There is an inherent tension between the Houston estimation and the long run changes in the model. The Houston estimation looks at a period up to nine years after the deregulation event. At this point, the new housing composition is still not equal to the stock, suggesting that the long run change in the stock of housing has still not been achieved. The natural question is: what would have happened after 2007?[7] Although new housing size remains small up to the end of the available data (around 2013), the analysis for that time period is not included in this paper for various reasons. First, there is notable volatility from differential shocks during the Great Recession; secondly, there were subsequent deregulation events in Houston after the Great Recession. These factors put into question any conclusions that can be made using the analysis after 2007.

Experiment	**Description**	**Long Run Δ Housing Size**
Baseline	Change in flow continues indefinitely	-14 log points
Alternative	Change in flow abruptly stops and matches stock	-3.8 log points

Without a credible way to get at causal estimates after 2007, I try to do two types of simulations that make assumptions about the long term nature of the short to medium term effect identified in Houston. The baseline simulation is to assume that the change in flow continues indefinitely; that is, the 14 log point drop in relative housing size will eventually cause the stock of housing to suffer an analogous 14 log point drop. In an alternative specification, I assume that the relative change in housing size of the stock (which is 3.8 log points by 2007) will cease to change any more; that is, the steady state will have been reached by 2007. Note that neither of these scenarios are particularly likely to be true, but they represent extreme scenarios that likely bound the true long term effect. I simulate these two effects to provide a sort of informal bound on the types of social welfare effects that the model outputs. As a final note, if I had to take a stance between the two outcomes, I note that economic intuition suggests that long term steady states do not tend to abruptly arise; hence, it is my view that the baseline simulation represents something closer to reality.

[7] The differential impacts of the Great Recession on housing supply and demand, and subsequent policy changes in Houston, make it difficult to make credible inferences past 2007.

To induce the model to decrease housing size in the long run by 14 log points (baseline) or 3.8 log points (alternative), the marginal price of housing size needs to rise by about 13 log points or 4.3 log points respectively. Since the estimated Houston drop in price was about 12.5 log points, this is further evidence that the reality of these housing markets may be more closely matched with the baseline simulation. The baseline model results are shown in the next section. In Appendix A.4.4, I show the results for the alternative specification, where the magnitudes of the effects are smaller, but the heterogeneity is still present.

1.4.3 Welfare Results

Figure 1.14: **Baseline Simulation:** Household Lifetime Gains Across Income and Household Size (2010 Dollars)

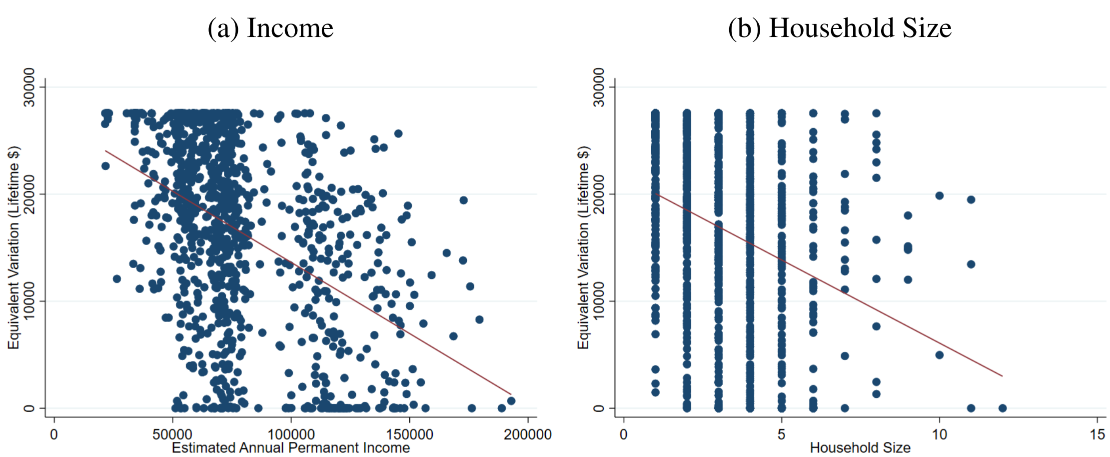

For the two scenarios, the model is simulated for two parallel worlds, one with the deregulation event and one without. Then, equivalent variation (the income needed to make a household indifferent between the policy change and the status quo) is calculated for each household. Note that some households will move into the city that deregulated. For the stayers in the deregulated city, the simulated deregulation event causes welfare gains throughout the income and family size distribution. This is not surprising, as the base price of housing is decreasing. Heterogeneity in welfare gains comes from their differences in elasticity of demand for housing size, which are regulated by their income and family characteristics. Note that this analysis looks at the rental value of housing consumption, and ignores the effects on asset prices for homeowners.

The welfare gains, for households who were always in Houston, are substantial. Over a lifetime, the deregulation amounts to $18,000 which is about one third of the median income. This is somewhat higher than the amount from the reduced form price regression in Table 1.5. More importantly, there is heterogeneity in lifetime gains across the income and household size distribution. Specifically, lower income and smaller households benefit more from the deregulation event. The equivalent variation varies by as much as $25000 between the households that have the least to gain (top decile in terms of household size and income) and those who have the most to gain (bottom decile).

1.4.3.1 Homeowners vs Renters

The previous section looks only at the rental value of consumption services. However, a calculation of interest also includes the effect on homeowners in terms of their housing prices. In Figure 1.15, I show the effects of the policy change but add in the shock to house values to homeowners as identified by the Census data. For many homeowners, their net gain is negative, meaning that any savings in rent and gains from the ability to adjust housing size is negated by the fall in their housing values. In many ways, this calculation is about the political feasibility of such a policy change. Because homeowners are more likely to be wealthy, including the implied asset price changes into their equivalent variation implies a range of heterogeneity that is one magnitude larger. Although the heterogeneity amongst homeowners is not as noticeable, the fact that there is strong selection into homeownership means that, on average, the slope across the income distribution is approximately the same as in the situation where everyone is a renter.

Figure 1.15: **Baseline Simulation With Asset Price Effects:** Household Lifetime Gains Across Income and Household Size (2010 Dollars)

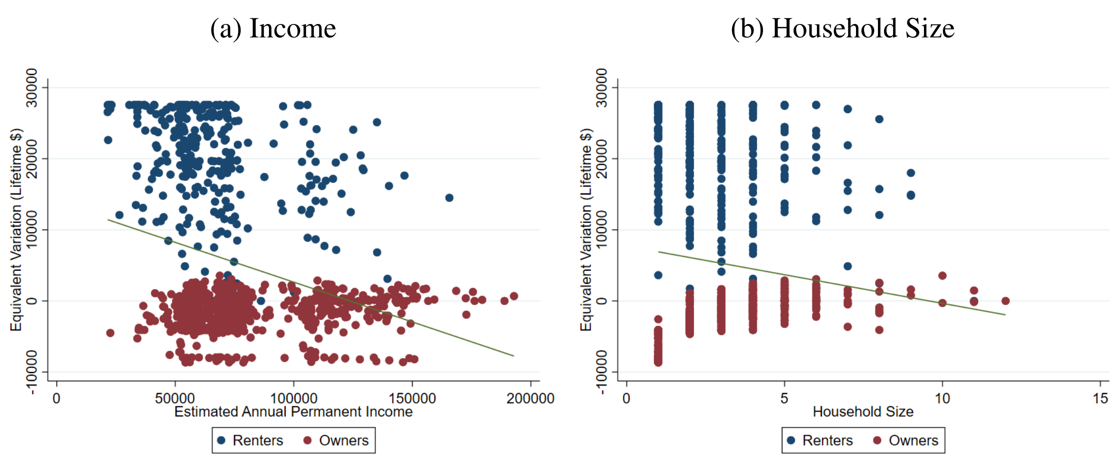

1.4.3.2 Welfare Decomposition: Lot Size Savings vs Re-optimization Gains

With the intuition derived in Section 1.2.1.1, the total welfare gains in Figure 1.14 can be decomposed into two parts: a part that comes from lot size savings conditional on initial house size demand, and a part that represents the gains that come from the ability to choose a new house size (re-optimization). I find that the first order effects are coming from lot size savings; re-optimization is economically important: ($2000 average gains). However, although re-optimization is increasing in income and household size, the heterogeneity is relatively small. Hence, the first order effects dominate.

Figure 1.16: **Baseline Simulation:** Value of Lot Size Savings Across Income and Household Size (2010 Dollars)

(a) Income (b) Household Size

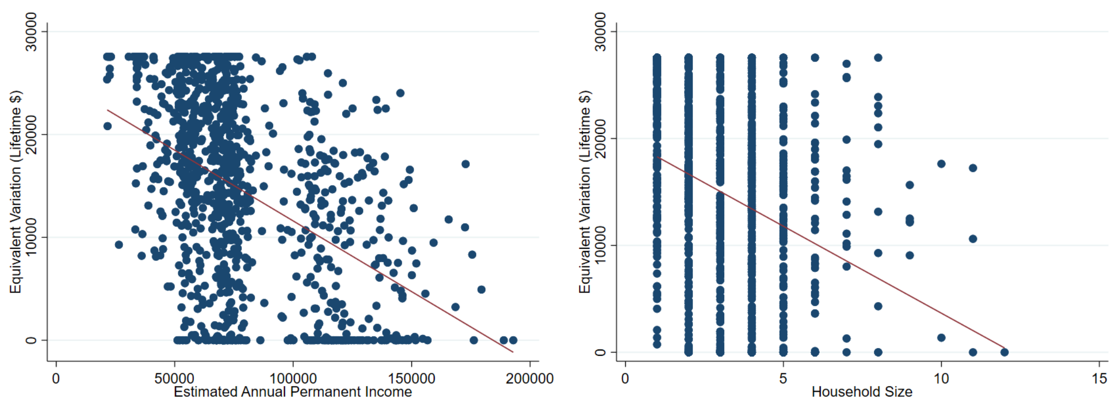

Notes: The value of lot size savings is simply the money saved given a household's initial house size demand, without the opportunity to re-optimize and choose a new house size.

Figure 1.17: **Baseline Simulation:** Re-optimization Gains Across Income and Household Size (2010 Dollars)

(a) Income (b) Household Size

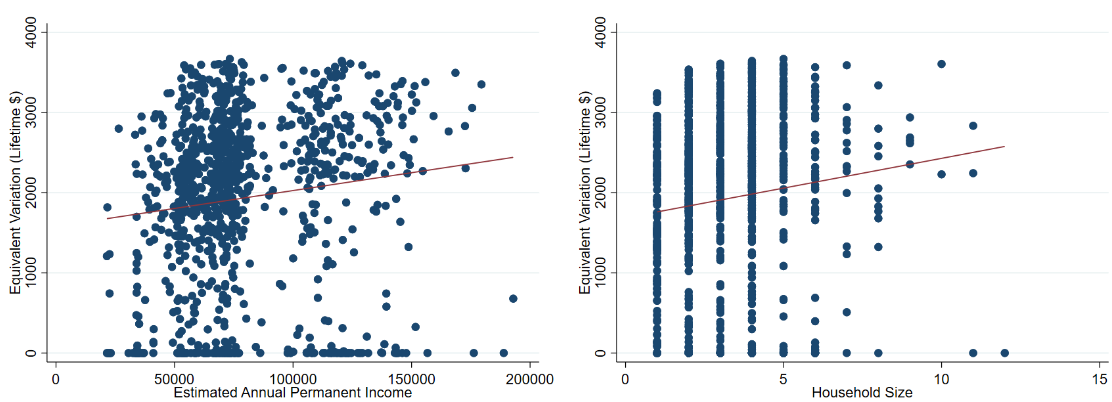

Notes: The value of re-optimization gains is the difference between total welfare gains and the initial lot size gains. It represents the value of being able to move to a smaller house in response to price changes.

1.4.4 Extension: Selection Results

In the previous section I show that the model is capable of rationalizing key features of the patterns of housing size and demography in the data. More importantly, it predicts important avenues of heterogeneity that are driven by demographic features but which work through the standard avenues of demand in response to changes in marginal prices. Another important prediction of this model, which I view as a success, is described in this section.

One of the main predictions of the model is the types of selection that would occur across cities, through the same mechanisms that select on housing size. Economic intuition implies that the Roy Model mechanisms will select on age, income, and HH size. The key inputs into the Roy Model are the correlations and variances between the locational preference shocks ζ_i^L. That is, by specifying the joint distribution of error terms, the selection of people into each city can be changed. To illustrate the model's tendency, I use uncorrelated error terms with the same variance; therefore, any type of resulting selection would arise from the correlations in utility generated by the demand problem (in terms of household size and income) and the underlying correlations in household size and income. For a further mathematical exposition of the Roy Model in this setting, see Appendix A.6 for the precise mechanisms for how such selection could arise. The theoretical result coming from this selection can be summarized as follows:

Proposition 1. *Under a first order log normal approximation of utility, if the correlation of household size and income is not too negative, lower income and smaller families will move into the deregulated city.*

In addition to the theoretical result, I show the simulated result from the model. Table 1.9 shows the equivalent difference-in-difference estimator for selection on characteristics as predicted by the model with normal uncorrelated locational preference shocks.[8] For simplicity, I show the results for the baseline model only. The differences are driven by the types of people moving into the affected city (Houston). As predicted, the affected city has smaller families and lower income through selection. In the model, this is driven entirely by this population's disproportionate gains from the deregulation of the MLS.

I then turn to the data from the Current Population Survey. Although the geographies available are not as precise as in the Corelogic dataset, I limit each city's sample to those people who live "in the central city". I run the analogous difference-in-difference specification where the pre-post periods align with the 1999 change in policy.

[8]The shocks were drawn from an uncorrelated normal distribution with a standard deviation of 200 for each city. The non-affected city had a higher mean draw of 500 in order to compensate households so that enough people would still live there after the welfare gain due to the simulated deregulation event. Note that the units here are in utility units, which do not have additional interpretable meaning. The magnitudes of these parameters are such that the resulting city sizes and effect sizes look reasonable. They do not affect the direction of selection.

Table 1.9: Selection of Deregulated City vs. Status Quo City

	(1) HH Size	(2) HH Income
$(\bar{X}_{2,post} - \bar{X}_{1,post}) - (\bar{X}_{2,pre} - \bar{X}_{1,pre})$	-0.186	-3988
$\bar{X}_{\text{2,post}}$	3.25	63353.4

Notes: Difference-in-difference estimates from simulation of baseline model (see Table 1.8) and simulated preference shocks (see footnote).

$$y_{ict} = \gamma_t + \lambda_c + \beta houston * post + \epsilon_{ict} \tag{1.6}$$

where y_{ict} is an outcome for household i in metro area c during year t. β thus represents the relevant difference-in-difference estimator that corresponds, theoretically, to the resulting direction of selection predicted in the model.

Table 1.10 shows the results from the basic difference-in-difference empirical model. As an extension, I also look at Age and College. There is little discernable effect on age, but Houston is relatively less college-educated, has smaller families, and has lower incomes. College education may be a better measure of permanent income than current income. For a naive policy maker, decreases in education and income are socially undesirable results, but from the perspective of selection in this model, it is merely a symptom of a minimum lot size decrease that actually favored certain demographics more than others; simply put, it indicates that the types of people moving into Houston are more likely to gain than others. In the context of a city (Houston) and state (Texas) that experienced relatively high levels of immigration (both internationally and from other parts of the United States), it may be realistic to assume that there is enough immigration for selection to be relevant, but more investigation into the breakdown of immigration and emigration flows may be required for further verification of this theory.

1.5 Conclusion

The general point of this paper is that MLS regulations have substantial effects on house size, prices, and household welfare. These effects are significantly heterogeneous. In my analysis, I take into account an additional channel that minimum lot size regulations operate through that the existing housing literature has ignored. MLS regulations increase house sizes by reducing the marginal cost of an additional square foot (i.e., "it is easier to build a big house on a big lot"). This intuition is verified empirically with the Houston event study, which shows that smaller minimum

Table 1.10: Selection of People into Houston vs. Other Texas Cities

	(1) Age	(2) College	(3) HH Size	(4) HH Income
Houston*Post	0.547	-0.059	-0.203	-4102.938
	(0.544)	(0.013)	(0.035)	(923.793)
Observations	18903	18903	18903	18903
Metro FE	x	x	x	x
Year FE	x	x	x	x
$\bar{X}_{\text{Houston,post}}$	46.2	0.291	2.71	63207.53

Notes: Data from IPUMS CPS. Sample consists of central city households in Texas from 1991-1997 (pre) and 1998-2007 (post) periods. Standard errors clustered by metro.

lot sizes led to smaller houses being built. Using a structural approach to model the features of housing size demand, I show that MLS regulations have large welfare costs, and these costs are unevenly distributed in the population. Specifically, it is families with fewer people as well as poorer people who are disproportionately hurt. To the extent that this demographic is younger, this is also potentially a generational issue.

These results came from plausible and reasonable considerations of all the nuances involved in housing demand. The analysis incorporated the nonlinear pricing of housing size given by changes in the minimum lot size regulation. The analysis also incorporated the necessity good (nonlinear Engel curve) features of housing consumption. The analysis also connected the price elasticities of house size demand from Houston to the model and showed they were consistent. These details are important in getting both the quantitative and qualitative results correct.

Note that this analysis looks at one aspect of housing regulations: its operational effect on the size of lots, which passes costs onto households. The heterogeneity of these costs is important, but they do not address the underlying reason for why these regulations exist in the first place. Such minimum lot size regulations may exist as exclusionary tools to solve freerider problems in public goods distribution. They may also correct for negative externalities of poor neighborhood and neighbor characteristics and increase the value of amenities. As such, this analysis is only a partial input into the full analysis that a social planner or policy maker might want to take into account. The analysis in the paper, in the context of all the effects of housing regulations, shines a spotlight onto the housing size channel that was previously left in the dark.

CHAPTER 2

The Demographic Determinants of Housing Regulations

2.1 Introduction

Housing regulations are very important determinants of local and aggregate economic outcomes. Therefore, understanding the development of these regulations gives policy makers information to implement policy interventions and gives researchers exogenous variation to study the causal effects of interventions. In the United States, restrictive regulations on residential land use are determined by a variety of regulations on the state and especially on the local level. Because these regulations are enacted locally, there exists widespread heterogeneity in the spatial distribution of these regulations (see Figure 2.1). However, the distribution of these regulations is not random; rather, they are highly correlated with local political alignments, incomes, and other variables. This paper investigates the role of the demographic determinants of housing regulations, particularly the role of fertility booms (most especially during the Baby Boom period between 1940 and 1960) in increasing housing regulations.

Figure 2.1: Distribution of Housing Regulations

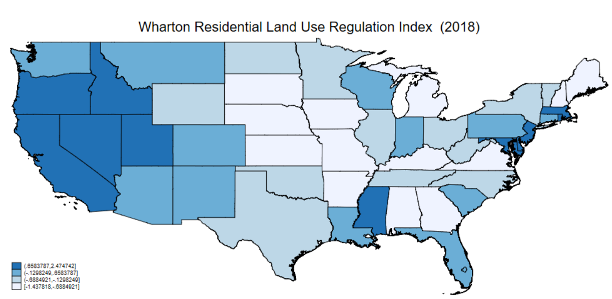

The theoretical motivation for this investigation is based on a logical extension of Fischel's Homevoter Hypothesis, described in Fischel (2005). The standard version is a theory of the political economy that posits that homeowners and their concern about property values motivate the regulation of land use. By regulating building so that supply is limited, this raises property values for existing homeowners. Because fertility booms may increase the demand for housing, any exogenous variation in fertility booms may be causally associated with increased demand for owning housing assets, and thus lead households to be more concerned about the value of these housing assets. The resulting political economy equilibrium may be one of more restrictive housing regulations.

One of the larger questions this paper wants to answer is how this aspect of the Homevoter Hypothesis interacts with increased demand for housing arising from demographic trends. By studying a critical historical period after the Second World War and the ensuing rapid suburbanization, one can make testable predictions about the patterns of housing ownership, fertility, and restrictive housing regulations in a disciplined way. Figure 2.2 shows the basic reduced form relationship between the development of housing regulations and the size of the Baby Boom. The focus on this period is important both because of the size of the Baby Boom but also because of the stickiness[1], or difficult-to-reverse nature, of these housing regulations which both restrict housing development but also (a) exclude certain types of voters and (b) change the incentives for those who move in. That is, when large fertility changes induce exclusive housing regulations, they exclude certain types of people who benefit more from deregulation and also create more incentive for existing voters to be more restrictive. In short, these fertility booms can induce localities to be "stuck" in high regulation equilibria.

Figure 2.2: Wharton Land Use Regulation Index (2005) and Size of the Baby Boom: State Level

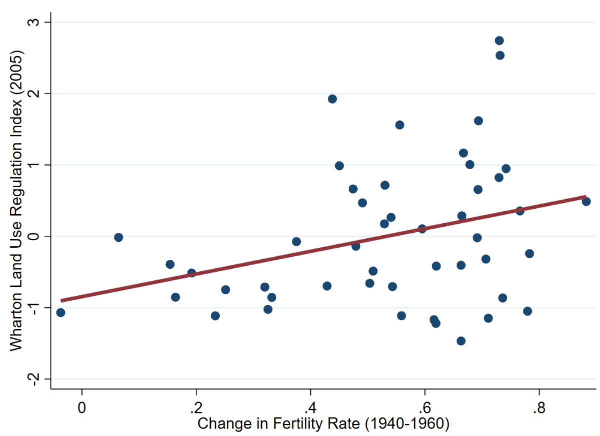

Note: Wharton Residential Land Use Regulation Index (2005) from Gyourko et al. (2008)

The mechanisms studied theoretically in this paper show how increased fertility leads to political economy equilibria with more housing regulations. Households trade off a benefit of owning

[1]The fact that housing regulations are almost never repealed is well-described in pages 77 to 79 of Levine (2005), who describes data from Massachusetts which shows that once single family zoning has been established, it is almost never relaxed.

housing against the risk of their housing assets depreciating. A fertility boom shifts this tradeoff towards the side of owning more housing, due to the larger needs of households in the consumption of housing. In a political economy equilibrium of the style of Ortalo-Magné and Prat (2014), the households get to vote on whether to issue new permits to build more housing in their local area. Hence, the median voter, in terms of their preferences and asset position, determines whether to increase housing supply. An increase in housing supply (i.e., permits to build) reduces rents but also reduces asset prices; hence, those with more exposure to asset price risk tend to be more opposed to increasing housing supply. Overall, exogenous increases in fertility lead to increases in homeownership and housing asset levels. This exposes households to more pain if house prices fall; hence, these households tend to vote for restrictive housing regulations that limit the size of their city.

The rest of the paper will go as follows: first, I highlight the literature and more thoroughly describe the historical period in which I analyze these trends. Second, I detail a political economy model where fertility shifts are added, and I look at the comparative statics of the political economy outcomes (notably city size) with respect to an exogenous increase in fertility. Finally, I document the fertility trends and housing regulation trends in the data and show they are consistent with the comparative statics in the model. I also use World War II mobilization rates as an exogenous shifter of fertility trends, to confirm a plausible causal connection between fertility rate changes and restrictive land use regulations.

2.2 Literature

The history of U.S. housing regulation, as described by Hirt (2015), Gyourko and Molloy (2015), and Whittemore (2020) involved a crucial period after World War II in which rapid suburbanization and changes in legal precedent incentivized many communities to enact restrictive housing regulations. These regulations varied in many forms, from minimum lot and floor space requirements to restrictions on the permit process of building. Quantitative measures of the number of these regulations like in Ganong and Shoag (2017) show that the quantity of these regulations increased steadily from WWII to the turn of the century. The stated motivations for these regulations included racial animus, conservation of neighborhood characteristics, and environmental concerns. Economic theory, as in Hamilton (1975) has also posited that these exclusionary housing regulations serve as decentralized solutions to the problem of excluding freeriders from public goods offered in their respective jurisdictions.

Fischel's Homevoter Hypothesis, outlined in Fischel (2005), sits prominently as one of the main explanatory frameworks for the political economy of housing regulations. Economic theorists, as in Ortalo-Magné and Prat (2014), have mathematically formalized the underlying mechanisms in

Fischel's broad theory, and derived intuitive economic mechanisms that illustrate the tension between homeowners' interests in the value of their housing assets and overall housing affordability.

There is an empirical literature that tests the Homevoter Hypothesis. The evidence for the hypothesis is mixed. Notably, cross-sectional comparisons in limited local areas, as in Dubin et al. (1992) or McDonald (1995), show that homeowners are more likely to support zoning ordinances. Closest to the empirical strategy in this paper is Hilber and Robert-Nicoud (2013), who uses the number of married people with no children in an area as an instrumental variable, but finds statistically insignificant effects. More importantly, there may be many directions of causality between fertility, marriage decisions, and political positions, so that instrument may not satisfy the necessary exclusion restriction.

This paper uses an instrument that borrows from the economic demography literature about the Baby Boom. That literature offers many explanations for the Baby Boom, from household technological innovation to improvements in maternal health, as well as behavioral explanations about relative deprivation during the Great Depression. One of the major explanations, however, remains the role of World War II, as explained in Acemoglu et al. (2004). Doepke et al. (2015) use the World War II measure of mobilization rates by state as an exogenous shifter of the fertility patterns across states. Goldin and Olivetti (2013) find nuance in WWII's effects on women's labor supply across educational attainment levels, suggesting that the direction of the effect may be ambiguous. This paper takes the latest iteration of this literature from Brodeur and Kattan (2022) who use novel county-level data and find a negative relationship between WWII mobilization and fertility. This instrument will be further described in the empirical section.

2.3 Contribution and Preview of Findings

The contribution of this paper is to (a) demonstrate the key theoretical assumptions needed to rationalize the key relationships between fertility rate changes, homeownership, and housing regulations in the political economy and housing markets (b) more credibly test the central mechanisms of the Homevoter Hypothesis through a new empirical approach involving exogenous shifts in fertility and (c) offer a credible alternative instrument for housing regulation indices for empirical economic research.

There are two main findings of this paper: one theoretical and one empirical. The first is that the positive correlations between fertility rate changes, homeownership, and housing regulations are a natural consequence of only a few additional assumptions added to the standard Ortalo-Magné and Prat (2014) political economy model. To this model I add an exogenously varying quantity of housing need (which is presumably affected by fertility) and diminishing marginal utility of consumption (while keeping the risk-averse preferences that provide a mean-variance tradeoff). The latter plausible assumption provides a natural positive relationship between fertility rate changes (which create changes in housing needs) and homeownership by increasing the marginal utility of homeownership. Lastly, I more thoroughly pin down the size of the city by adding congestion externalities of new residents; this is a parsimonious way to rule out multiple equilibria and more specifically pin down a simple positive relationship between homeownership and housing regulations.

The second main finding is that the relationship between the size of the Baby Boom and housing regulations is quite robust; these relationships hold at both the state level and using newer county-level data, and with a variety of covariates as controls. Moreover, using plausibly exogenous variation in World War II mobilization rates that affected fertility, I show that the causation mechanisms by which fertility could have increased the development of housing regulations have economically and statistically significant magnitudes. This is an interesting fact as a matter of historical economic research, but it may also be relevant to a researcher who is looking to find a suitable instrument for housing regulations. Using WWII induced fertility booms to instrument for housing regulations might be an improvement or a compelling alternative to existing ways in which housing regulations are instrumented for, as in Saiz (2010) and Quigley et al. (2008) who use the local public expenditure share or measures of religiosity to instrument for local housing regulations.

2.4 Model

The model captures a few key features of the causal mechanism. First, larger families face larger total rental costs because they need more housing quantity (for example, in the form of space); this is the primary mechanism. Second, an increase in the size of the family should lead to more incentive to own, rather than rent, housing. Finally, an increase in family size (from a Baby Boom, for example) should lead to the median voter favoring a more restrictive permitting process (a smaller city size, for example).

In other words, are the theorized demographic mechanisms here consistent with the Homevoter Hypothesis? Specifically, is the Ortalo-Magné and Prat (2014) model capable of matching the empirical patterns? There are several innovations I make on top of that standard model: first, the use of quadratic utility instead of standard mean-variance utility gives diminishing returns to wealth in expectation; this means that changes in housing needs shifts marginal utility in a way that affects housing ownership decisions. This channel is key to explaining the positive relationship between fertility and homeownership. Secondly, I characterize, in this new context, a maximum size of the city which reflects more housing regulations, as well as the comparative statics of how that size changes with respect to a change in housing needs.

2.4.1 Model Environment and Timing

This is an overlapping generations model simplified from Ortalo-Magné and Prat (2014) so that each generation only lives two periods. The income process and the idea of a political economy equilibrium are similar, but I use a different utility function with differing housing needs (fertility) and I pin down the size of the city by introducing a convex congestion cost function.

There are two locations: the countryside and the city. The city has space for some mass N that is less than or equal to 1, and the city's expansion (perhaps via the use of permits) is voted on by the households. The countryside can accommodate everyone. The outside option of the countryside is standardized so utility is some level u_0.

The households consume their wealth w_e at the end of their life. They maximize the expectation of a quadratic utility function:

$$E[U(w_e)] = E[-(w_e - \bar{w})^2] + \epsilon_i \tag{2.1}$$

where $\bar{w}$ is a parameter which is sufficiently large so that the utility function is strictly increasing in w_e and ϵ_i is a preference parameter for living in the city.

Table 2.1: Timing of Decisions

Time	**Events for one generation (in order)**
t-1	HH Born, knows current income Buys housing amount h Votes to expand city or not for next period
t	New city income is realized Rent is paid, House is sold Consumption of wealth HH Dies

In the city, income is the random walk process:

$$y_t = y_{t-1} + \mu_t \tag{2.2}$$

where μ_t is an i.i.d. variable with variance σ^2.

There are financial markets where households can borrow and lend freely at the exogenous risk-free rate $\frac{1}{\beta}$. The timing of events and decisions are outlined in Table 2.1, and the budget constraint follows this timing. The explanation for voting is left to the end.

Given the timing of the events, end of life wealth for the relevant cohort at time t can be written as:

$$w_e = y_t - \lambda r_t + p_t h_{t-1} - p_{t-1} h_{t-1} \frac{1}{\beta} + r_t h_{t-1} \tag{2.3}$$

where

1. λ is an exogenous, family-size parameter that determines the scale of rents and income[2]

2. r_t is the rental rate per unit quantity of housing

3. p_t is the asset price per unit quantity of housing

4. h_{t-1} is the quantity of housing purchased at the beginning of life.

Essentially, the household's balance at the end of their life is their realization of their household income (y_t), minus their rents (λr_t). If they choose to buy some quantity of housing h_{t-1} in the previous period, they encounter the user cost of that housing ($p_t - p_{t-1}\frac{1}{\beta}$), but save $r_t h_{t-1}$, which is the proportional part of their rent.

The owners of rental housing exist outside of the city; they are risk-neutral and can borrow and lend in the same financial markets as the household. However, there is a wedge between rental prices r_t and the flows to the owners of rental housing. Specifically, the owners of rental housing only realize flows of $r_t - \rho$ where ρ represents some relative disadvantage in flow revenue. As in Poterba and Sinai (2008), ρ can represent tax costs or some additional maintenance costs. This wedge drives the arbitrage-free price of houses down. The result of this is that households find it less costly to own housing than to rent housing.

2.4.2 Prices and Equilibrium

In such models, it is known that in equilibrium, rent r_t is a linear function of city income (with an additional constant $\bar{r}$). I first start by "guessing" with this pricing function. Then I check that it is an equilibrium. $\bar{r}$ is a constant in the price function that adjusts to clear the housing market by inducing more or fewer people to live in the city as opposed to the countryside.

$$r_t = \frac{y_t}{\lambda} + \bar{r} \tag{2.4}$$

Standard asset pricing, combined with the price wedge ρ implies that

$$p_t = \frac{\beta}{1-\beta}\left(\frac{y_t}{\lambda} + \bar{r} - \rho\right) \tag{2.5}$$

[2]An alternative specification involves scaling rents and income. The mechanisms in that model would be the same. There many also be microfoundations for this type of rental costs.

The reason that households do not all own the full amount of housing is because they face the risk of asset depreciation from the aggregate city income shocks. Because they commit to buying housing before the city realizes its aggregate shock, the household trades off the benefits of owning housing against the disutility of risk. This can be seen by an equivalent maximization problem that is consistent with the utility maximization of the given quadratic utility function.

The problem can be written as:

$$\max_{h_{t-1}} E[U(w_e)] = \max_{h_{t-1}} \left[-(E[w_e] - \bar{w})^2 - Var[w_e] \right] \tag{2.6}$$

Essentially, the household is trading off some increasing quadratic function of expected wealth against the disutility coming from the variance of that wealth.

Taking prices as given, the expectation and variance of wealth are the following:

$$E[w_e] = -\lambda\bar{r} + \rho h_{t-1} \tag{2.7}$$

$$V[w_e] = \frac{h_{t-1}^2 \sigma^2}{\lambda^2(1-\beta)^2} \tag{2.8}$$

In Appendix B.1, I show that optimal total housing consumption is:

$$h_{t-1}^* = \frac{\rho(\lambda\bar{r} + \bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}} \tag{2.9}$$

Homeownership, in partial equilibrium, is therefore increasing in family size λ. The primary effect, in the numerator, is the larger need for housing rents. The larger rents effectively shifts wealth consumption in a direction that increases marginal utility, and thus increases the incentive to own rather than to rent housing; its effect parallel an increase in the fixed part of rent $\bar{r}$ or the consumption point $\bar{w}$. The equilibrium pricing function also scales volatility to clear the market, so there is an additional effect coming from reducing the variance term in the utility function; this effect shows up in the denominator. However, this effect is secondary and not the focus of this model.

In Appendix B.2, I show that expected utility is given as below:

$$E[U] = \underbrace{-\Big(1 - \frac{\rho^2}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\Big)^2 (\lambda\bar{r} + \bar{w})^2 - \Big(\frac{\rho(\lambda\bar{r} + \bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\Big)^2 \frac{\sigma^2}{\lambda^2(1-\beta)^2}}_{F(\bar{r})} + \epsilon_i \tag{2.10}$$

This function is strictly decreasing in $\bar{r}$. This is not surprising as utility should be decreasing as the fixed part of rent $\bar{r}$ increases. Given parameters, write the utility function above as a function of $\bar{r}$. Call the part of the function without the preference draw ϵ_i as $F(\bar{r})$. Since this function is strictly decreasing in the relevant area of wealth, there exists a well-defined inverse.

For a household to live in the city, it must be that their utility in the city is higher than the outside option u_0. That is:

$$F(\bar{r}) + \epsilon_i > u_0 \tag{2.11}$$

The fixed part of the rental price $\bar{r}$ in the pricing function is pinned down by the marginal household in the city. To see this, assume ϵ_i, the preference parameter for living in the city, has a distribution characterized by some cumulative distribution function $Z(e)$. For some person with a preference draw of $\epsilon = \epsilon^*$, market clearing means this must hold:

$$Z(\epsilon^*) + N = 1 \tag{2.12}$$

For housing supply to equal housing demand in the city, the person who has the preference draw ϵ^* must be indifferent between living in the city and living in the countryside. Hence, there must be a fixed point where

$$F(\bar{r}^*) + \epsilon^* = \mu_0 \tag{2.13}$$

where $\bar{r}^* = F^{-1}(\mu_0 - Z^{-1}(N-1))$ since the inverse is well-defined. The derivation of $\bar{r}^*$ is given in Appendix B.3. In other words, the fixed part of rent is determined by the surplus of living in the city by the marginal household.

2.4.3 Political Economy Equilibrium

I have shown the equilibrium conditions and prices that clear the housing market when N is given. However, in the model, households get to vote on whether to expand N in the next period, which necessarily changes prices.

$$d\bar{r} = \frac{\partial F^{-1}(\mu_0 - Z^{-1}(N-1))}{\partial N} dN \tag{2.14}$$

Because the original function F is strictly decreasing, this derivative must be negative. That is, an increase in the size of the city decreases rents and asset prices, specifically the fixed part $\bar{r}$. This is intuitive. But how does the household weigh the costs and benefits of whether to increase N?

The effect on the utility of the household consists of multiple parts. First, the resulting decrease in rent increases their utility. Second, the resulting decrease in asset prices lowers their utility. Third, there is a marginal *congestion cost of expansion* $f'(N)$ that each households pays; this cost can represent a combination of physical congestion costs, psychic costs of crowding, infrastructure costs of building, public goods costs, permit revenue and other things. What is important is that these costs are low (or negative, in the case of permit revenue) when the city is small and very big (e.g., approaching infinity) when the city is large. This, in essence, pins down the preferred size of the city as households vote on whether each marginal increase in city size is worth it. This cost reflects the fact that existing large cities that push against their boundaries have high costs of expansion.

As a whole, what determines whether a change in N is beneficial depends on the asset position h that has already been decided relative to the population of the city. That is, the change in wealth dW_e can be decomposed as follows:

$$dW_e = \underbrace{-\lambda d\bar{r}}_{\text{marginal benefit: savings on rent}} + \underbrace{\frac{1}{1-\beta}\hat{h}_{t-1}d\bar{r} - f'(N)dN}_{\text{marginal costs: fall in assets and congestion costs}} \tag{2.15}$$

where $d\bar{r}$ is defined as above. The first term comes from rental costs and the second term comes from asset price changes in addition to the net congestion costs of expansion.

Thus, setting Equation 2.15 to zero represents a political economy equilibrium, for which the household will no longer vote to expand the city. This is the point right before when the marginal cost of expansion is too high to justify the benefits of saving on rent. This political economy equilibrium is stable because the costs of expansion are exponentially increasing to infinity, and hence dominate the resulting tradeoffs generated between renting and keeping house prices high.

Note that the tradeoff between renting and housing prices is an example of a classic game theoretic failure that illustrates the main points of the Homevoter Hypothesis. Because voting is done by a generation of voters who already made their housing asset decisions, having a high level of these assets (which gives these households more exposure to asset price changes) makes them support a lower city size than is otherwise optimal. This is also a consequence of there being no mechanism for future households to provide transfers to compensate current generations for their losses due to city expansion.

2.4.3.1 Effect of Housing Needs Shock

What, then, is the effect of an exogenous increase in λ? To illustrate this conceptually, I plot the marginal benefit and marginal cost curves in Figure 2.3. The intersection between marginal cost and marginal benefit denotes the original political economy equilibrium with a city size of N_{old}^{PE}. Suppose λ_{low} increases to λ_{high}. This has two effects: first, it increases the marginal benefit because every unit of rent savings goes farther for the household when the household consumes more housing overall. This, on the surface, would make households support a larger city size. However, under the standard interior solution assumptions, the marginal cost increases dominate the marginal benefit increases, and as a result, the resulting city size in a political equilibrium decreases to N_{new}^{PE}.

Figure 2.3: Marginal Costs and Benefits of Voting for City Size

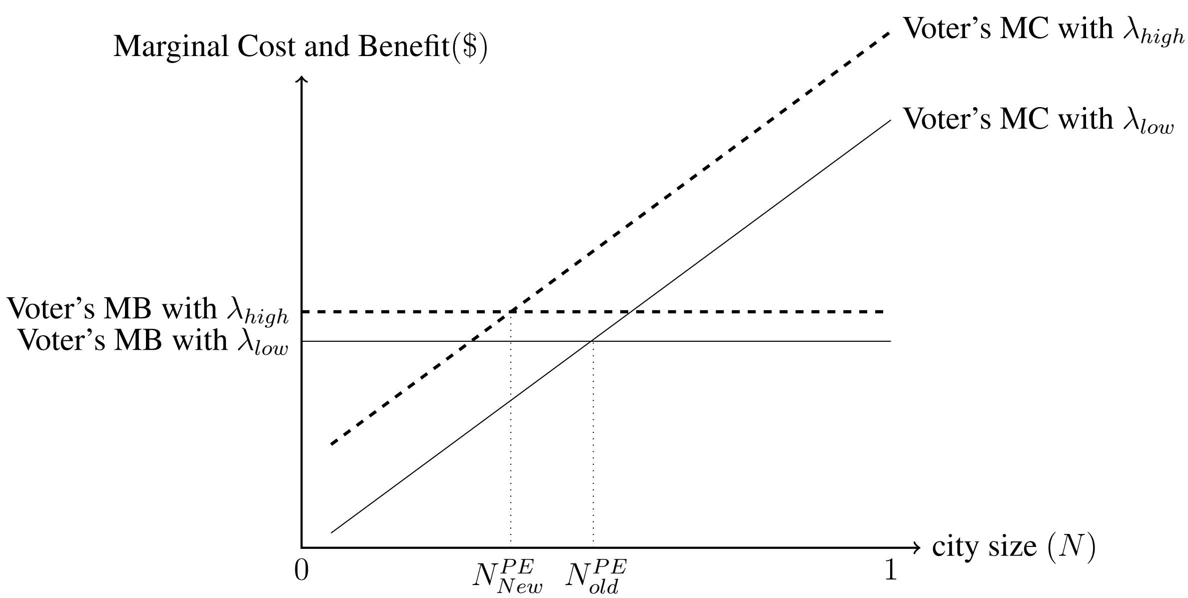

In short, it can be shown that the following result holds:

Proposition 2. *Under some mild regularity conditions for an interior solution, N^{PE} is decreasing in λ.*

This is the precise statement of what housing regulations increasing with fertility means in this model: fertility increases cause higher housing needs, and the mechanisms translate that to allowing fewer permits for building in the voting process. In Appendix B.4, I parameterize and simulate a version of this model that illustrates the results above quantitatively.

There are several intuitions for this result. The first is that asset prices reflect perpetuity values, so any changes in homeownership is magnified by the multiplier $\frac{1}{1-\beta}$. The second is that homeownership responds to housing need (fertility) changes directly through the channel of decreasing marginal utility as described before. Note that marginal benefit also directly increases due to higher housing needs. The intuition for why the marginal costs shifts are larger is because changes in asset values (which represent some perpetuity value) tend to dominate the one-time savings on rent.

2.4.4 Model Summary

This is an example of a plausible model which illustrates a particular way in which a fertility boom could induce the development of housing regulations. The mechanisms involved are intuitive. It is not possible to exhaustively rule out alternative mechanisms by which these causal chains happen, but the empirical work in the following sections show that the underlying patterns of fertility, homeownership, and local land use regulations are consistent with the model.

2.5 Data

2.5.0.1 Housing Regulation Data

The residential land use regulation data comes from two main sources. The first is the established Wharton Residential Land Use Regulation index for both the 2005 and 2018 years. This measure is described in Gyourko et al. (2008) and is an index measure of survey responses that ask jurisdictions a variety of questions about their regulations on land use for residential purposes. The second is the ranking of states according to their number of court cases per capita related to land use disputes, the main measure of regulation provided by Ganong and Shoag (2017). This measure is available for both 2005 and 1975. Their paper also uses a separate 1975 survey of jurisdictions for robustness, which I also use. This measure is essentially a measure of disputes; it is a function of both the presence of restrictive laws as well as the presence of parties willing to fight over those laws, perhaps due to higher demands for building. In total, I have five total measures of land use regulation at three different years. For comparison purposes, I standardize each variable by demeaning and rescaling to standard deviations. For all five measures, a higher value represents more restrictive housing regulations. In Appendix B.5 I map the geographical distribution of each measure.

2.5.0.2 Fertility Data

Two measures of fertility are used. The first follows Doepke et al.'s methods by taking public census data (IPUMS) from the 1940 and 1960 census and measuring the average number of children age 5 or under, per woman. This is aggregated at the state level. The second measure of fertility is the birth/population rate at the county level from Brodeur and Kattan (2022).

2.5.0.3 WWII Mobilization Data

The idea of using WWII mobilization rate was pioneered by Acemoglu et al. (2004), but mobilization rates are only available at the state level. Since then, Brodeur and Kattan (2022) has compiled county level data for WWII casualties, which are highly correlated with mobilization rates. As such, I use county level WWII casualty rate data as a proxy for WWII mobilization.

To use WWII mobilization rates as an exogenous instrument, there are several concerns to address. Mobilization rates differed based on race, education, and local economic conditions. First, the U.S. military remained relatively exclusionary for people of color. Second, college education was a possible alternative to military service. Third, draft exceptions were made for people working in agricultural industries (in order to secure the country's food supply). As such, the instrument is conditional on the local Black share, the college educated share, and the farmer share. Conditional on these variables, the assumption is that the remaining variation consists of plausibly exogenous variation from the lottery nature of the draft and the idiosyncracies of the local draft board.

2.5.0.4 State and MSA Level Analysis

The baseline sample of the MSA level regressions consists of the largest 280 MSAs that exist today (but its size is measured by the 1940 population). I also restrict the sample to the largest 135 MSAs for robustness in Appendix B.6. All of the MSA level analysis is done with region fixed effects. Controls for education, income, black share, and farmer share are added. These controls are consistent with the reasoning in Brodeur and Kattan (2022), where instrument validity will require conditioning on variables on which there was likely selection into WWII mobilization.

2.5.0.5 Summary Statistics

Table 2.2: Summary Statistics: State Level

	Mean	SD	Min	Max	N
Fertility Rate (1940)	1.89	0.30	1.41	2.77	48
Fertility Rate (1960)	2.41	0.27	1.35	3.15	48
Δ Fertility Rate (1940-1960)	0.53	0.21	-0.04	0.88	48
WWII Casualty Rate (%)	0.24	0.04	0.16	0.38	48

Table 2.3: Summary Statistics: MSA Level

	Mean	SD	Min	Max	N
Δ Fertility Rate (1940-1960)	0.64	0.41	-1.76	2.50	280
WWII Casualty Rate (%)	0.24	0.05	0.13	0.52	280
Average Income	772	208	214	1271	280
Average Education	11.4	1.18	6.66	14.1	280
Percentage Black	9.9	14.0	0	64.7	280
Percentage Farmer	16.7	13.5	0	61.8	280

2.6 Empirical Work

First, I establish some relevant correlations that are consistent with the theory. The main correlation is that changes in fertility rates should be positively correlated with housing regulations.

2.6.0.1 Timing of Baby Boom vs Baby Bust

The focus is on the 1940-1960 change in the fertility rate because of several reasons. The main reason is that the effects on housing regulation seem to be strongest for the time periods in which Baby Booms are strongest. Over larger periods of time that included both the Baby Boom and Baby Bust (say, between 1940 and 1970) the relationship between housing regulations and fertility is close to zero correlation and not statistically significant. The second reason is that this period is special in terms of the direction and magnitudes of the effects; every state during this period had a positive fertility shock, and there is more variation in terms of the size of those shocks. Thus, this period, which is a critical period of increasing fertility, seems to be most relevant for inducing equilibria where households may be stuck in a high-regulation, low-size city. In Table 2.4, I show the historical relationship between the two main measures of housing regulations available after 1975 and the size of the fertility changes in each state. Only in the Baby Boom period (1940-1960) do I find a positive correlation between fertility and more housing regulations.

2.6.1 Positive Relationship Between Housing Regulations and Fertility - State and MSA Level Analysis

The basic quantitative relationships between fertility rate changes and measures of housing regulation are expected to be positive. I find that they are positive on the state level and the MSA level. However, the relationship is stronger for earlier surveys of housing regulations. Later measures of housing regulations (as in the WRLURI 2018) are less correlated with the size of the Baby

Table 2.4: State-Level Housing Regulations and Historical Fertility Changes

	(1) WRLURI (2005)	(2) WRLURI (2005)	(3) GS (1975)	(4) GS (1975)
Δ Fertility Rate (1940-1960)	2.457 (0.914)		2.943 (1.084)	
Δ Fertility Rate (1940-1970)				0.119 (1.192)
Δ Fertility Rate (1940-1990)		0.452 (1.425)		
Observations	48	48	48	48
Income + Education Controls	x	x	x	x
Race Controls	x	x	x	x
Fertility Rate (1940)	x	x	x	x

Boom. It is difficult to test exactly if there is evidence of a fading relationship between WWII mobilization and housing regulations today, but this is at least consistent with historical narratives where housing regulations reflect historical changes during the critical period after WWII. I focus on the 2005 WRLURI. Below I run the basic ordinary least squares regression model at the state and MSA level.

For the state level, I use the full vector of state-level controls. The relationship between the size of the fertility boom and the level of housing regulations is positive. Table 2.5 shows that this relationship is persistent across different measures of housing regulations, and is statistically significant for most measures, even while controlling for racial and income factors.

$$R_s = \lambda \Delta f_s + \beta X_s + \epsilon_s$$

where R_s is a housing regulation measure for state s, Δf_s is the change in each state's measured fertility between 1940 and 1960. X_s are a vector of controls.

Table 2.5: State-Level Housing Regulations

	(1) WRLURI (2005)	(2) WRLURI (2018)	(3) GS (1975)	(4) GS (2005)	(5) Survey (1975)
Δ Fertility Rate (1940-1960)	0.944 (0.645)	0.649 (0.732)	1.951 (0.691)	1.878 (0.574)	1.835 (0.798)
Observations	48	48	48	48	48
Income + Education Controls	x	x	x	x	x
Race Controls	x	x	x	x	x
Income Growth Controls	x	x	x	x	x
Fertility Rate (1940)	x	x	x	x	x

For the MSA level regressions, I regress the WRLURI (2005) Index on the 1940-1960 changes in fertility, all with region effects. For the OLS regressions, I show that the inclusion of additional control variables matters. In fact, the relationship between housing regulations and the size of the Baby Boom in Table 2.6 is not statistically significant with the full vector of controls; this is consistent with the idea that fertility rates are not fully exogenous. For example, if the development of religiosity or political conservatism increases fertility but also reduce housing regulations, then this omitted variable would lead to a downward bias of the OLS estimate.

Table 2.6: MSA-Level Housing Regulations Regressions

Sample: Largest 280 MSAs

	(1) WRLURI (2005)	(2) WRLURI (2005)	(3) WRLURI (2005)
Δ Fertility Rate (1940-1960)	0.410	0.212	0.195
	(0.191)	(0.188)	(0.191)
Observations	280	280	280
Average Education		x	x
Average Income (1940)		x	x
Black Percentage			x
Farmer Percentage			x
Region FE	x	x	x

2.6.2 Positive Relationship Between Fertility and Homeownership - State Level and HH Level Analysis

The next empirical question is whether fertility rates increase homeownership as suggested in the model. The results show there is a strong positive relationship between fertility and homeownership, even after controlling for income, race, location, and other variables. States that had the largest fertility booms had significantly higher increases in homeownership.

I also run a more detailed regression with the cross-section of the 1960 Census. Specifically, the full specification is:

$$h_i = \gamma_l + \zeta C_i + \theta X_i + \xi_i \tag{2.16}$$

where h_i is a measure for homeownership for households i, γ_l denotes location fixed effects, C_i is the number of children under five. X_i is a vector of controls which includes income and age, squares of income and age, and other demographic variables. To measure homeownership, I use the standard homeownership variable, which is a binary variable.

Table 2.7: Homeownership and Fertility: State Level Changes

	(1) Δ Homeownership (1940-1960)	(2) Δ Homeownership (1940-1960)	(3) Δ Homeownership (1940-1960)
Δ Fertility Rate (1940-1960)	0.234	0.211	0.225
	(0.061)	(0.060)	(0.063)
Average Income (1940)		0.000	
		(0.000)	
Pct White (1940)		-0.112	
		(0.061)	
Δ Income (1940-1960)			0.000
			(0.000)
Δ Pct White (1940-1960)			0.088
			(0.142)
Constant	0.203	0.258	0.136
	(0.026)	(0.052)	(0.058)
Observations	48	48	48

Table 2.8: Homeownership and Fertility: Cross-Sectional HH Regressions

	(1) Homeownership	(2) Homeownership	(3) Homeownership	(4) Homeownership
Children Under 5	0.020	0.023	0.022	0.017
	(0.000)	(0.000)	(0.004)	(0.003)
Family Members Above 5	0.042	0.040	0.037	0.033
	(0.000)	(0.000)	(0.005)	(0.003)
Observations	2642347	2642347	2642347	2642347
Income + Age Controls	x	x	x	x
Race + Sex Controls		x	x	x
State Fixed Effects			x	
County Fixed Effects				x

The results in Tables 2.7 and 2.8 show that, even controlling for a long list of demographic and location controls, homeownership is increasing in fertility both in the difference specification (by state), as well as the cross-sectional specification. All of this is consistent with the basic patterns predicted by theory.

2.6.3 WWII Mobilization Rate: Exogenous Variation in Fertility

The empirical analysis in the previous section looks at changes in fertility rates as they are; however, many models of fertility have joint decisions that are related to household income. Moreover, political alignments can influence both attitudes about fertility and attitudes toward restrictive housing regulations.

To get around the problem of the endogeneity of fertility, I use a proxy for WWII mobilization (WWII casualty rate) published by Brodeur and Kattan (2022). The fundamental idea is that areas

(MSAs in this context) which sent more men to fight in WWII had relative increases in the labor supply of women. The introduction of these women into the workplace may have made it easier for subsequent women to enter the workforce; hence, in places that had more mobilization, the women had fewer children because of the higher opportunity costs of having children.[3]

2.6.3.1 First Stage

Table 2.9: MSA-Level First Stage Regressions

	(1) Δ Fertility Rate (1940-1960)	(2) Δ Fertility Rate (1940-1960)	(3) Δ Fertility Rate (1940-1960)
WWII Casualty Rate (%)	-2.249 (0.441)	-2.368 (0.431)	-2.360 (0.480)
Observations	280	280	280
Average Education		x	x
Average Income (1940)		x	x
Black Percentage			x
Farmer Percentage			x
Region FE	x	x	x
F-Stat (Excluded Instrument)	26	30.1	24.2

As shown in Table 2.9, running the regression of the change in fertility rates on the WWII casualty rate gives a strong negative relationship with a slope of -2.4 and a t-statistic of about -4.9. The F-Stats are also reported, and they range from 24 to 30 depending on the vector of controls. There is coefficient stability with the estimates; the underlying relationship between WWII casualty rate and fertility rate does not change very much across the different specifications.

2.6.3.2 Reduced Form

The first specification of interest is the direct, reduced form relationship between the exogenous variable and the outcome of interest. In Table 2.10 I find that a 1 percentage point increase in the WWII casualty rate (which is a significant increase) is associated with a 3.6 to 4 standard deviation lower level of housing regulations as measured by the WRLURI in 2005.

[3] Brodeur and Kattan (2022) also looked at the direct effect of deaths on overall fertility, but they conclude that WWII death rates were not significant enough to cause significant variation in fertility.

Table 2.10: Relationship Between Wharton Residential Land Use Regulation Index (2005) and WWII Casualty Rates

Sample: Largest 280 MSAs

	(1) WRLURI (2005)	(2) WRLURI (2005)	(3) WRLURI (2005)
WWII Casualty Rate (%)	-3.684	-3.911	-3.984
	(1.454)	(1.393)	(1.555)
Observations	280	280	280
Average Education		x	x
Average Income (1940)		x	x
Black Percentage			x
Farmer Percentage			x
Region FE	x	x	x

2.6.3.3 IV - 2SLS

Lastly, I specify the instrumental variables estimation via the Two-Stage Least Squares specification below:

$$R_i = \theta\widehat{\Delta f_s} + \eta X_i + \epsilon_i$$
$$\widehat{\Delta f_s} = \delta Z_i + \psi X_i + \mu_i$$

I instrument for the change in fertility rates between 1940 and 1960 with the WWII mobilization rate. The identifying assumption is, conditional on controlling for income, education, Black share, and farmer share, that mobilization rates affect housing regulations only through the channel of fertility. The results from Table 2.11 should be compared to the OLS results in Table 2.6. In fact, the magnitude of the estimated effects are much higher, although these estimates are much more noisy with the significantly larger standard errors. These IV results show that looking at the variation in fertility that is plausibly exogenous gives results that are consistent with the theory: increases in fertility rates cause the development of more housing regulations that persist at least until 2005.

2.7 Conclusion

In conclusion, I have demonstrated that a simple political economy model, with modest modifications, can rationalize the empirical patterns that we see in fertility, homeownership, housing assets, and housing regulations. Specifically, these additional assumptions involve exogenous shifts in housing demand and declining marginal utility of wealth, which build on top of the mean-variance

Table 2.11: 2SLS Estimates: Housing Regulations and Changes in Fertility (Instrumented with WWII Casualty Rates)

Sample: Largest 280 MSAs

	(1) WRLURI (2005)	(2) WRLURI (2005)	(3) WRLURI (2005)
Δ Fertility Rate (1940-1960)	1.757 (0.651)	1.652 (0.648)	1.688 (0.725)
Observations	280	280	280
Average Education		x	x
Average Income (1940)		x	x
Black Percentage			x
Farmer Percentage			x
Region FE	x	x	x

tradeoff in the Ortalo-Magné and Prat (2014) political economy model. The addition of an increasing congestion cost of expansion pins down the city size in a political economy equilibrium by ruling out multiple equilibria. This historical perspective is important for understanding how housing regulations came to be and the mechanisms by which welfare losses happen due to the regulation. In summary, exogenous changes in housing needs due to fertility changes ultimately lead to more restrictive housing regulations (smaller city size) that hurt average welfare.

To further investigate this historical narrative, I use state-of-the-art WWII mobilization data, available at the county level from Brodeur and Kattan (2022), as an instrument. I explain its plausibility as an exogenous source of variation for fertility. After conditioning on a variety of control variables that likely affected selection into WWII military service, I show that the instrument is relevant for the size of the Baby Boom. The results are generally consistent with the idea that WWII mobilization decreased fertility rates. The instrumental variables estimation results are generally consistent with the idea that higher fertility rates cause more housing regulations.

CHAPTER 3

Housing Vacancy Chains

3.1 Introduction

There is considerable political and public policy debate about the distribution of gains from building market price (unsubsidized) private housing. Often, these developments are criticized as unaffordable and catering to the upper classes. To the extent that markets for new housing are segmented from the overall housing market, the development of this housing may provide supply only to a relatively exclusive group of people. There is an alternative view, grounded in a theory about the nature of the housing quality ladder, that suggests building new housing on the higher end has spillover effects that help people who cannot afford or are otherwise not interested in higher end housing. These spillover effects exist because of housing vacancy chains: building new housing creates vacancies by having households vacate their existing house and create additional vacancies, a process which continues on. If these vacancy chains dip into varied socieconomic groups or neighborhoods, then there are potentially large spillover effects of building new housing.

To help answer this question, I use real estate purchase transactions data to trace the migration chain of owners who move. This is accomplished using fuzzy name-matching of buyer/seller names with deed transactions data. These transactions are merged with public mortgage application data to deduce characteristics about the movers. This project characterizes the level of segmentation in the network of buyers and sellers, which is then used to infer a causal effect of building a new housing unit on different demographic groups.

This paper is the first contemporary study of U.S. housing chains that identifies individual or household level characteristics. It also borrows heavily from the intergenerational income literature to characterize the extent of segmentation and correlation amongst buyers and sellers. One main caveat is that this data is limited only to owner occupied housing within a given area, which precludes directly observing the churning that arises due to renters or outsiders entering a particular metropolitan area under study. However, the finding that, even amongst real estate purchase transactions, interlink elasticities are low (that is, high socioeconomic status for sellers is only weakly correlated with high socioeconomic status for buyers) suggests that not looking at renters would only understate the extent of spillover effects of building new housing.

The study of vacancy chains should be distinguished from a related concept, which is the effects of filtering. Although definitions differ depending on the context, filtering is the aging of new housing into less expensive older housing. This concept of depreciation may be empirically important, but its effects take years, and even decades, to realize. Moreover, the downward effects on housing prices through the aging process may be easily reversed by substantial investments in housing improvements (renovations), as well as neighborhood-wide regulations like historical preservation districts which both limit the supply of housing and maintain a high level of amenities. In contrast, the study of vacancy chains is the study of economic processes that move on the order of days, weeks, or months; the study is therefore more temporally relevant for the incumbents of a commuting zone or an interconnected set of jurisdictions.

In general, I find a broad series of evidence that there are widespread and immediate spillover effects of new housing construction. At the core of my findings is a relatively low interlink elasticity of characteristics like income and sales price, and a relatively low persistence of characteristics like race along the links of these vacancy chains. This is means that houses that were built and sold disproportionate to a particular racial group or income group create vacancies for comparatively lower socioeconomic groups several links down the vacancy chain. I find that following the household level characteristics along these links on the chains, rather than census-tract characteristics, can be quantitatively important to get accurate estimates because existing studies using census-tract level characteristics tends to underestimate such elasticities. Overall, these findings are largely consistent with existing studies studying vacancy chains in other contexts or with less precise data.

3.2 Literature Review

There is a limited literature on the use of vacancy chains to study the effects of new housing built. Lansing et al. (1984) physically track housing migration chains in the United States during the 1960's. More recent work has used proprietary or administrative data to track these chains, notably Mast (2021) and Bratu et al. (2023). Generally, these studies find that the chains that start from new housing builds reach lower income groups rather quickly.

A related literature covers heterogenous movers or segmented housing search. Landvoigt et al. (2015) study the San Diego housing market with different types of movers. Piazzesi et al. (2020) document that broad searchers across different types of housing markets help these markets respond to supply or demand shocks. This paper is also inspired by matching algorithms, both for matching sellers to buyers across links in the transactions data, as well as matching the transactions data to mortgage application data (HMDA). In particular, Anenberg and Ringo (2022) and Rosen (2011) use similar algorithms.

This paper also uses measures of interlink elasticities and rank-rank slopes that are inspired by the literature on intergenerational income, specifically from Chetty et al. (2014) and Deutscher and Mazumder (2021). Finally, this paper connects to a broader set of research in the fields of housing policy and urban planning, with White (1971), Turner (2008) and Ferrari (2011) studying the theory and empirical evidence of housing vacancy chains in different contexts, often relying on Markov assumptions with transition matrices.

This study of vacancy chains also intersects with a lot of existing research on how housing is segregated among racial, class, and other lines. For this paper, I focus first on income (both tract and household level income), but analyze a variety of other outcomes including house values, neighborhood characteristics, racial group, and loan type. The reasons for different types of segregation are varied. A canonical explanation is that individual preferences lead to sorting, as in Schelling (1971) and Becker and Murphy (2003). Shertzer et al. (2022) argue that public policies like zoning play both a historical and contemporary role. Hardman and Ioannides (2004) show that even though there is some evidence of urban mixing, there exist pockets of segregated housing in geographical space. Massey et al. (1987) argues that such segregation may be costly for household well-being, especially for the Black population.

This paper is the first to study household and individual level of movers in the contemporary United States context. Existing studies that use proprietary change of address data lack individual level characteristics, but instead ascribe households or individuals the characteristics of their neighborhood. This paper offers nuances about how using tract-level characteristics can be slightly misleading, but does not find evidence that overturns the major results of existing research on housing vacancy chains.

3.3 Empirical Analysis

3.3.1 Data Description

The first dataset consists of the Corelogic deed / transactions data, which is a collection of administrative data available in local jurisdictions for recording real estate transactions. The relevant variables include buyer and seller names, transaction prices, transaction dates, new construction indicator, mortgage amounts, and lender name. Moreover, each property is precisely geocoded.

The second dataset is the publicly available Home Mortgage Disclosure Act (HMDA) data available from the Consumer Financial Protection Bureau. The dataset consists of the universe of mortgage applications. This dataset has information that Corelogic does not have, like household reported income, as well as individual race and gender variables.

The sample consists of residential real estate transactions within the Los Angeles metropolitan area, including all transactions from 2005 to 2016, with identified new house construction restricted to be after 2007.

3.3.2 Empirical Strategy

There are two major parts to the empirical strategy. The first strategy looks at every housing transaction in the sample from 2010-2016 and tries to match it with characteristics of a relevant buyer who sold their house (possibly going back before 2010). By tracing where the buyer sold a house around the same time they bought a house, I can infer what types of neighborhoods they are moving from, the sales price of their previous house, as well as household specific characteristics like income and race if they applied for a mortgage. This simple strategy of tracing one link back is called the **One-Link Analysis**, which looks at the correlation between one step in the seller and buyer characteristics regardless of whether the houses are new or not. Part of the identifying assumption is that the structural correlation between sellers and buyer characteristics are invariant to whether the house being transacted is old or new, and also invariant to unobservables like how many links separate the (possibly old house transaction) from the newly built house transaction. That is, if there's nothing special about new housing other than being expensive or catering to the wealthy, then the results of this analysis should be accurate.

The second strategy is simply an extension of the first strategy, which is to continue to look for buyer and seller matches to trace out the vacancy chain that *start with newly built housing*. This is the **Vacancy Chain Analysis**, which conditions the analysis on the subset where the full chain is traced. Because there is attrition of the sample due to statistical mismatch and outsiders (renters included) buying housing, the number of matches during each step of the chain gets dramatically reduced. Hence, there is a tradeoff between these two approaches. This vacancy chain approach solves the problem of unobservables related to being closely connected with newly built housing, but may suffer from other sorts of selection bias based on matching at each step.

Below I detail the general algorithm for how the links are matched, and how all of these transactions are matched to HMDA, the mortgage application data.

3.3.3 Construction of Housing Transactions Links and HMDA Merge

The construction of vacancy chains follows the following algorithm which is largely based on Anenberg and Ringo (2022):

1. Generate an initial set of properties (either all of them or those associated with a new construction flag)

2. Given each buyer's name, find a seller who has an exact match on the last name, fuzzy match on the first name (which includes the middle initial), and has a transaction date within a year of the previous transaction. The fuzzy match requires the two strings to be within a given Jaro-Winkler distance.[1]

3. If the above match is not unique, look for a match where the other listed buyer/seller first names have a fuzzy match. Prioritize missing second buyer/seller names over strong non-matches of second buyer/seller names.

4. If the above match is still not unique, keep the transaction that is closest in date to the previous transaction. If there is still a tie, randomly choose one.

5. For the vacancy chain (multi-link analysis): Given the buyers in the next link of the chain, repeat the above step to find the next set of buyers in another link in the chain. Repeat until the sample is still sufficiently large.

The merging of real estate transactions with HMDA is done with the following algorithm:

1. Generate all potential matches that have the same census tract and year.

2. Those transactions that match uniquely and exactly on loan amount and lender names are considered matched.

3. Those transactions that match within $1000 of loan amount and match on lender names are considered matched.

3.3.4 One-link Analysis: Summary Statistics of Links and Merges

Matching rates and counts are given in each step. For the 2010-2016 sample of housing transactions in the Los Angeles metropolitan area, a total of 491,066 transactions were identified.

Table 3.1: Buyer/Seller Matching Statistics

Link	Total Count	Percentage of Houses Linked
Link 0	491066	100%
Link 1	69979	14.2%

Note: Link 0 represents all 2010-2016 transactions in the sample.

[1]See Winkler (1990) for details about how this distance measure works. The default cutoff for the Jaro-Winkler distance is 0.03. Robustness results for alternative values of this cutoff are given in Section 3.3.10.

Matching rates for links are about 14%. The non-matches would include cases of renters buying property, outside movers buying into the Los Angeles metropolitan area, common name combinations that are not unique, and true matching errors.

3.3.4.1 Summary Statistics

Table 3.2: Summary Statistics for Various Samples

Variable	Link 0 (Full Sample)	Link 0 (Matched to Link 1)	Link 1
Minority Population (mean)	55.1%	52.8%	56.7%
Median HH Income (mean)	73374	76577	75512
Poverty Rate (mean)	10.6%	9.9%	10.1%
Sale Amount	460300	501877	490205
Observations	491066	69979	69979

Note: Means for first four variables are observed at the census tract level. These values are equivalent to weighted averages of census tract characteristics, where weights are determined by observations of residential housing units in the Corelogic dataset.

Overall, the matched 0 link (the sellers) are different from the overall population of sellers in the sense that they live in neighborhoods that are higher income, less minority, lower poverty, and have more expensive housing. The next link features houses that are closer to the mean in terms of their characteristics.

3.3.4.2 Transaction Dates of Match

To show the characteristics of the link across time, I plot the difference between the transaction dates (buying and selling) that were linked to a purchase of the same person / household. For a positive amount X on this distribution, that represents a person who sold a house X days after they purchased a house. These graphs show that there is a bulk of mass close to zero for all links, suggesting that there is coordination to buy and sell a house at approximately the same time.

3.3.4.3 Interlink Elasticities

The approach to understanding the impact of new housing down the housing vacancy chain is to measure the association of income and sales amount characteristics along the chain. Some key measures, taken from the literature on intergenerational income measures, include the intergenerational income elasticity and the rank-rank slope[2].

[2] Robustness checks with rank-rank slope are reported in Appendix C.1. I find qualitatively the same results as the estimates of elasticities.

Figure 3.1: Link 1 - Date Gap Distribution: Sell Date Relative to House Purchase Date

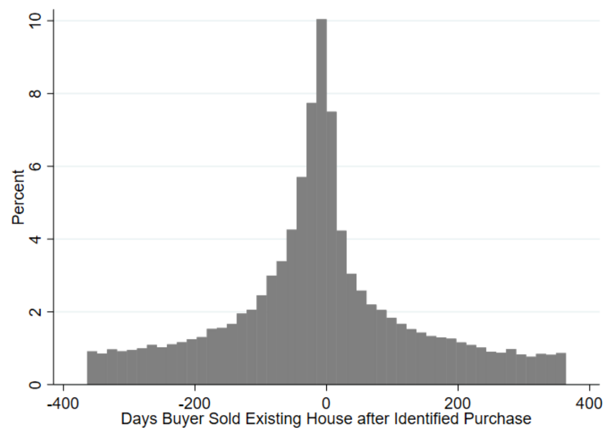

The interlink elasticities measure the strength of the relationship between seller characteristics and buyer characteristics across each subsequent link. Specifically, it is the slope coefficient β from the following regression:

$$\log(Y_{i,buyer}) = \alpha + \beta \log(Y_{i,seller}) + \epsilon_i$$

where c represents a distinct transaction. The outcome variable represents the characteristics of the buyer who bought the house that was sold by a seller with characteristics $Y_{i,seller}$. A value of β close to 1 represents a perfect correlation of a particular characteristic, and a value of β close to 0 represents no correlation (e.g., completely random house search behavior).

3.3.4.4 One-Link Elasticity Estimates: Income and Sales Amount

Table 3.3: Interlink Elasticties of Income and Sales Amount

Dependent Variable: Characteristics of Buyer

	(1) Log Median Tract Income	(2) Log Median Tract Income	(3) Log Household Income	(4) Sale Amount	(5) Sale Amount
Log Median Tract Income (Seller)	0.296 (0.004)	0.313 (0.007)			
Log Household Income (Seller)			0.364 (0.007)		
Sale Amount (Seller)				0.443 (0.004)	0.517 (0.007)
Constant	7.851 (0.040)	7.686 (0.075)	7.329 (0.078)	7.168 (0.051)	6.222 (0.092)
Observations	69973	18891	18891	69979	18891
Sample	Full Sample	Linked HMDA	Linked HMDA	Full Sample	Linked HMDA

Standard errors in parentheses

I find that the elasticities of income, measured on a census tract or individual level, range from 0.30 to 0.36. These values differ for many reasons: the first is the selection of households who were matched to HMDA, but the difference between the first two regression specifications suggest that these differences may be small. Elasticities estimated using HMDA data (household level), however, tend to be higher, perhaps reflecting strong correlation of unobservables that are not measured simply with census tract-level data. While the difference between 0.28 and 0.34 can be economically significant, I consider them relatively low income elasticities.

To put these estimates into more context, I use the thought experiment similar to the one described in Deutscher and Mazumder (2021) to calculate how many links needed for an outlier to regress near to the mean of the sample in income, given the estimated elasticities. A 0.3 elasticity of income implies that a house bought by a household who has a log difference of 1 (about 2.71x average income) should take, on average, about 3 links to reach a household within 0.03 log difference (about 3% difference) of average sample income. Similarly, if the elasticity is 0.36, then it would take 4 links on average. Because these links are theoretically occurring as quickly as weeks and months, this suggests that vacancies are quickly being created for a significant bulk of households who are average income or below.

The elasticities of sales amount represent the correlational strength between the sales amount of the house that the buyer sold for before buying the seller's house. That is, if one sells their house for $100,000, and bought another house for $100,000, that would represent datapoint that indicates a strong correlational relationship between those two sales amounts. The estimates suggest that the sales amount elasticities are much higher than the elasticities for income. For the Linked HMDA sample, it is 0.52. Running a similar thought experiment as above, it would take 6 links for a house 2.71x more expensive than average to reach a house within about 3% of the average sales amount. The higher correlation for this variable suggests a very important lesson about vacancy chains: while there is wide heteregeneity of impacts of housing vacancies across income, the sales amounts tend to be much more correlated across housing chain links. This may be because of credit constraints and time-invariant preferences for housing consumption. For example, if one's house could be sold for a particular amount, then financing restrictions might limit a future house purchase to nearly that same amount. In addition, preferences for types of housing, size of housing, and quality of housing means that subsequent sales and purchases may look similar in terms of sales amount.

3.3.4.5 One-Link Elasticity Estimates: Tract Level Poverty and Minority Percentages

I find similar qualitative results looking at the interlink elasticities of tract-level poverty and minority percentages. Specifically, the elasticities for the tract-level minority percentage is about 0.43 and the elasticities for tract level poverty percentage is about 0.23. These elasticities seem to be

Table 3.4: Interlink Elasticties of Neighborhood Poverty and Racial Characteristics

	Dependent Variable: Characteristics of Buyer			
	(1) Log Tract Minority Pct	(2) Log Tract Minority Pct	(3) Log Tract Poverty Pct	(4) Log Tract Poverty Pct
Log Tract Minority Pct (Seller)	0.425	0.444		
	(0.004)	(0.007)		
Log Tract Poverty Pct (Seller)			0.234	0.235
			(0.004)	(0.007)
Constant	2.293	2.217	1.528	1.471
	(0.015)	(0.026)	(0.008)	(0.015)
Observations	69979	18891	69614	18806
Sample	Full Sample	Linked HMDA	Full Sample	Linked HMDA

Standard errors in parentheses

above and below the elasticities of income, respectively. This suggests that racial composition may matter a lot more than poverty percentage in house search behavior, a finding consistent with the literature on housing segregation and racial preferences. However, an elasticity of 0.43 means that a house built in a very low minority neighborhood would, in expectation, create vacancies for a neighborhood within 3% of the average minority percentage in only 5 links into the vacancy chain.

3.3.4.6 One-Link Race and Loan Type Transition Matrix Estimates

An alternative way to understand categorical variables is to estimate transition matrices. Like the previous analysis with regressions of an autoregressive form, this approach has the standard Markov assumption that housing search behavior after a link are conditionally independent. That is, if Household 1 buys a house from Household 2 who buys a house from Household 3, then Household 1 is only affected by the characteristics of the owners and the house associated with Household 2; after conditioning on Household 2's characteristics, Household 1's search behavior should all be independent of Household 3 and their housing characteristics. Under these Markov assumptions, estimating a transition matrix and iterating them through time gives the unconditional long run characteristics of the population. For this analysis, I focus on two categorical variables using the HMDA-matched sample: the exact racial group and the type of loan.

Racial and ethnic groups are combined to be one of four categories: Non-Hispanic White, Hispanic White, Black, and Other Minority (which includes Asians, Pacific Islanders, Native Americans, and other groups). The loan type is either a conventional loan or an FHA/VA loan. This analysis tells whether there is persistence in these characteristics as transactions are made from buyer to seller. If the diagonal elements of the matrix are sufficiently high relative to the population percentages, then houses that are built for one racial group or one type of financing group may have impacts that are primarily within that group, rather than elsewhere.

Table 3.5: One-Link Transition Matrix: Race

Buyer Race	Seller Race: Non-Hispanic White	Hispanic White	Black	Other Minority
Non-Hispanic White	0.5548	0.2819	0.4125	0.4069
Hispanic White	0.1607	0.4609	0.2527	0.1766
Black	0.0261	0.0359	0.0913	0.0249
Other Minority	0.2584	0.2213	0.2435	0.3917

The transition matrix for seller/buyer race in Table 3.5 shows that there is persistence in this characteristic. The first matrix entry shows, for example, that of sellers who bought another house who are non-Hispanic whites, about 55% of the buyers of these sellers' houses were also non-Hispanic whites. The percentage below the first entry shows that about 16.1% of the houses that were owned by non-Hispanic whites were sold to Hispanics. This all shows that persistence exists but is not too strong. A significant fraction of such houses are still offered to buyers who are minorities.

Table 3.6: One-Link Transition Matrix: Loan Type

Buyer Loan Type	Seller Loan Type: Conventional	FHA/VA
Conventional	0.6977	0.4817
FHA/VA	0.3023	0.5183

Similarly, the transition matrix for loan type in Table 3.6 shows that there is persistence in this characteristic. One may be concerned about the persistence of this characteristic because FHA and VA loans are targeted government loan programs for certain populations that are deemed more vulnerable or in need of help. This exercise has several public policy implications: namely, building new housing of any type can have spillover effects on the other group, and new housing need not be built for FHA or VA populations specifically for there to be subsequent vacancy effects.

To illustrate the implications of these transition matrices and the point of links, I start with a fraction of new houses that are first owned by non-Hispanic whites or first conventionally financed. Under the standard Markov assumptions for transition matrices, I show the percentage of ownership by racial group and loan type expected in each link on the housing vacancy chain, and compare them to both the long run implied steady state (at infinity) and the observed population percentages in the sample. For different racial and ethnic groups, by building new housing, the steady state is

Table 3.7: One-Link Analysis (multi-link transition matrix): Race Impulse Response

Link	0	1	2	3	...	∞	Actual
Non-Hispanic White	100%	55%	47%	45%		44%	46%
Hispanic White	0%	16%	22%	23%		24%	23%
Black	0%	3%	3%	3%		3%	3%
Other Minority	0%	25%	29%	29%		29%	28%

nearly reached only about three links in. The percentages differ no larger than 1 percentage point from the implied steady state. This suggests that the spillover effects of building new housing, even if they are concentrated in one racial group, are rather broad. A similar results holds for loan type, where conventional financing of loans eventually reaches close to the steady state population just three links into the chain. Finally, I compare the steady states implied by the transition matrices to the observed population percentages in the HMDA-matched sample. The percentages are within a few percentage points of each other, but actual percentages are slightly higher for non-Hispanic whites and conventional mortgage loans. This is all consistent with the fact that the housing market at any given time is characterized by new housing builds: to the extent that these new builds are concentrated within those higher socioeconomic groups, we expect population percentages to be higher than steady state percentages.

Table 3.8: One-Link Analysis (link-specific transition matrices): Race Impulse Response

Link	0	1	2	3	...	∞	Actual
Conventional	100%	70%	63%	62%		61%	64%
FHA/VA	0%	30%	37%	38%		39%	36%

In conclusion, the One-Link Analysis reveals quite a lot about the relationships between buyers and sellers and the ways in which vacancies might propagate through in housing markets. I find that after a reasonable number of links (3-6), it will be difficult to tell what the characteristics of the original newly built house (the house that would have originated the chain) are. I do find some evidence that using household level income (instead of tract-level income) leads to higher interlink elasticity estimates, but these elasticities are not so high as to fundamentally undermine the conclusion that building new housing might have strong spillover effects into different socioeconomic groups.

3.3.5 Vacancy Chain Analysis: Summary Statistics of Links and Merges

This next section takes the analysis in the previous section and goes several steps further. Instead of looking at the relationship between buyers and sellers (one link), this analysis starts with new housing built and tries to trace the housing vacancy links back across the chain. This requires fewer assumptions about the Markov nature or autoregressive structure of vacancy chains because the nature of these chains are not inferred from studying the entire sample, but limited to the actual multi-link chains traced out by name matching. For clarity, I define the sequential nature of links below.

Definition of links: Link 0 represents the characteristics of the new house. Because these chains trace vacancies (in an abstract order, not necessarily in time), I label the house that was sold to buy the new house **Link 1**. Likewise, the next link on the chain (the house that was sold to buy the Link 1 house) is labeled **Link 2** and so on.

3.3.5.1 Vacancy Chain Analysis: Matching Rates

Matching rates and counts are given in each step. For the 2007-2016 sample of new housing construction in the Los Angeles metropolitan area, a total of 116789 transactions of new housing construction were identified.

Table 3.9: Buyer/Seller Matching Statistics

Link	Total Count	Percentage of Houses Linked
Link 0	116789	100%
Link 1	24766	21.2 %
Link 2	7036	28.4%
Link 3	2262	32.1%

Note: Link 0 represents new housing construction.

3.3.5.2 Vacancy Chain Analysis: Summary Statistics

Table 3.10 compares the characteristics of housing transactions between the different subsets. The old housing represents all the transactions that were for non-new construction, although a particular house transacted may not be old. Comparing the first column to the second column shows that new housing is built in neighborhoods with higher income and lower poverty rates, but new housing may actually be built in neighborhoods with a higher minority population. Moreover, there is a substantial difference in the sale amount. Overall, this is consistent with new construction being different from existing homes in terms of the characteristics of the buyers, as well as where the housing is built. Restricting the new construction to the chain subset, and looking only at the new construction (Link 0), I find that the subsample is remarkable similar to the full sample of new construction, suggesting that there are no strong selection effects when looking only at the identified chains. However, there may be a difference in the sale amount, but the differences between the third column and the second column seems to be a lot smaller than the differences between the first and the second.

Table 3.10: Summary Statistics for New Construction vs Existing Housing

Variable	Existing Housing (Full Sample)	New Construction (Full Sample)	New Construction (Link 0)
Minority Population (mean)	57.54%	59.9%	59.7%
Median HH Income (mean)	71139	75566	76510
Poverty Rate (mean)	11.3	10.8%	10.5%
Sale Amount	438713	505965	531204
Observations	1572061	116789	2262

Note: Means for first four variables are observed at the census tract level. These values are equivalent to weighted averages of census tract characteristics, where weights are determined by observations of residential housing units in the Corelogic dataset.

3.3.5.3 Vacancy Chain Analysis: Transaction Dates of Match

Similar to the transaction date analysis before, I show the transaction date difference between each link in Appendix C.2. The results are qualitatively very similar, with a large bulk of mass near zero. This suggests there is coordination of transaction dates between selling one's own house and buying another house.

3.3.6 Vacancy Chain Analysis: Tract and Transaction Level Average Characteristics

The following are the main results of the buyer/seller vacancy chains as constructed. The variables of interest are the median income, minority percentage, and poverty percentage of the tract, as well as the total sales amount of the transaction.

The median tract income figure shows that the subsequent households who move into housing created by vacancies in these chains tend come from, on average, lower income tracts compared to the households that bought the new house. However the variance is relatively large, with some chain links reaching to census tracts with median income as much as 2.3x compared to the median tract income of the new construction at Link 0.

The distribution of sales amount in the subsequent links is actually rather close to the distribution of sales amount for the new houses. This is also consistent with the fact that the interlink elasticities for sales amount is relatively high.

Figure 3.2: Median Income (Tract) by Link

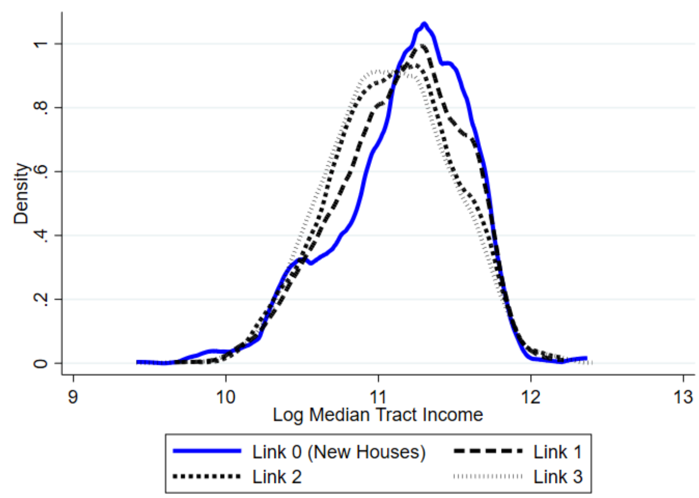

Figure 3.3: Sales Amount by Link

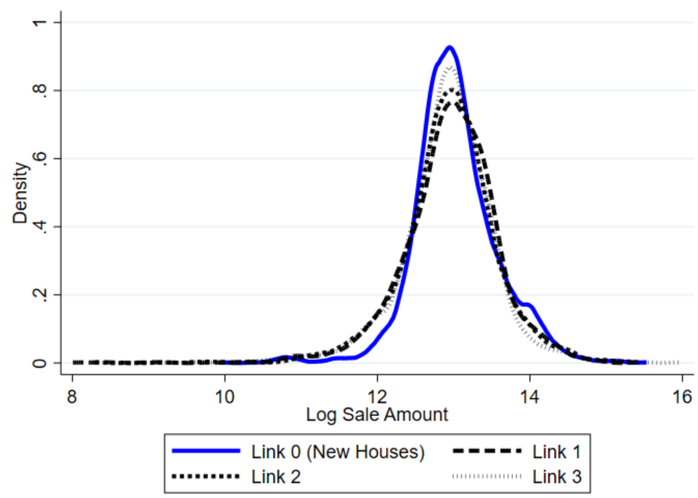

The minority population percentage results show a remarkable change in the distribution of where these houses are located as one analyzes the difference chains. Notably, new house construction is likely to be in neighborhoods with lower minority percentage than links further into the chain. Compared to new construction, Link 3 features houses in neighborhoods that are about 10 percentage points more minority on average. A similar result holds qualitiatively for poverty percentage, but the median and average do not change too much. However, the tail of the distribution is changing: there are more houses that are being transacted in lower income neighborhoods by link 3.

Finally, alternative visualizations of this data is provided in Appendix C.4; the results are qual-

Figure 3.4: Minority Population (Tract) by Link

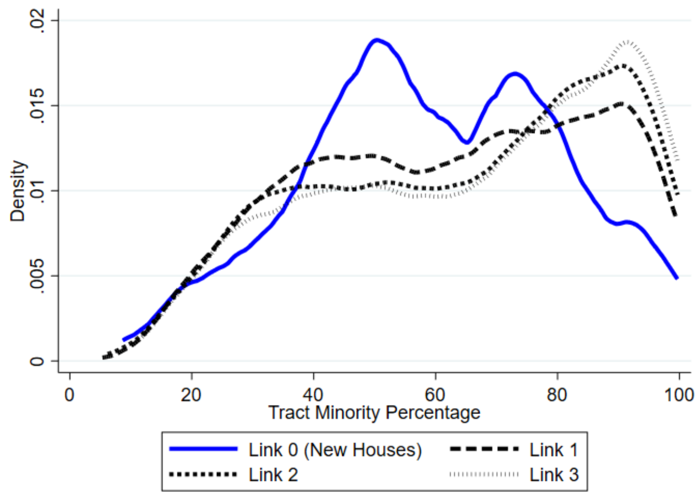

Figure 3.5: Poverty Percentage (Tract) by Link

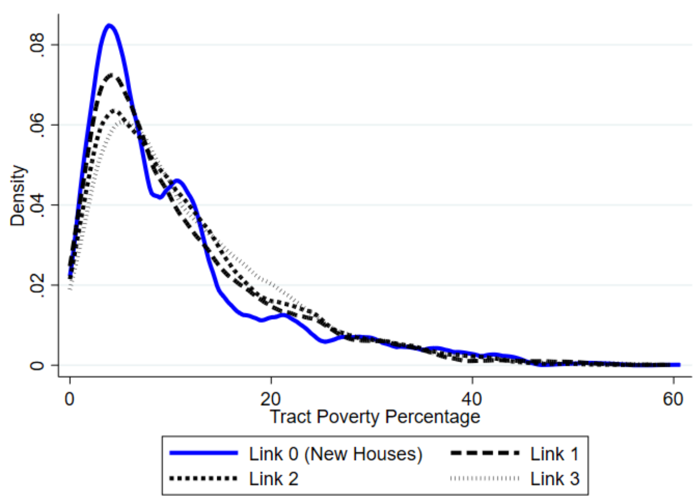

itatively similar.

3.3.7 Vacancy Chain Analysis: HMDA Match Rates

To further characterize the heterogeneity of homeowners across these housing vacancy chains, I merge the identified transactions in each link with the HMDA mortgage application data, as detailed in a previous section. The percentage matched to HMDA is not perfect; but there are enough observations to analyze a changing distribution. Moreover, match rates are similar across different links.

Table 3.11: HMDA - Corelogic Merge Statistics

Link	Total Transactions	Percent Merged w HMDA
Link 0	2262	36.4%
Link 1	2262	39.6%
Link 2	2262	37.2%
Link 3	2262	36.8%

Note: Link 0 represents new housing construction.

3.3.8 Vacancy Chain Analysis: Household Level Results (HMDA-matched sample)

Figure 3.6: Log Household Income by Link

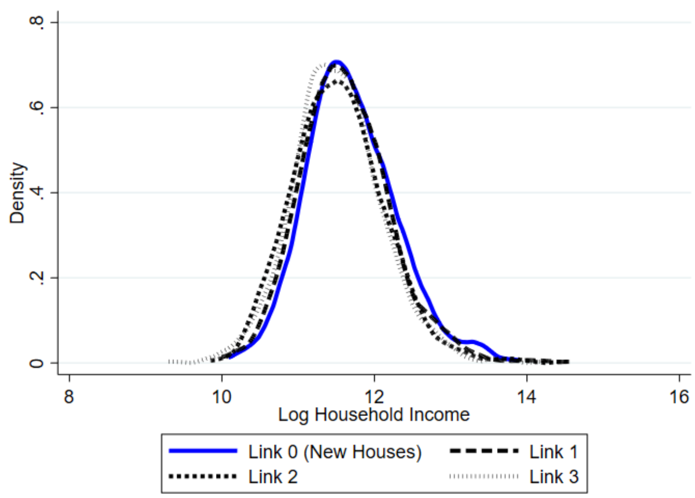

The results in Figure 3.6 show the distribution of reported household income (on mortgage applications) for each link along the vacancy chains. There is some evidence of a leftward shift of the distribution along each link from Link 0 to Link 3, suggesting that average income is regressing towards the mean. Compared to Figure 3.2, these shifts are less noticeable.[3] This is consistent with the idea that there exists within-tract variation in income that is very important. That is, tracking transactions using tract-level income may tend to overstate the amount of variation there is in actual household income; this result is also consistent with the higher elasticity of income estimates obtained using household level income instead of tract level income.

[3]A standard Kolmogorov–Smirnov test at the 1% significance level, for both of these figures, shows that these distributions are statistically distinct between Link 0 and Link 3.

3.3.9 Vacancy Chain Analysis

I use the same types of analysis of elasticities and transition matrices in the One-Link Analysis on the identified Vacancy Chains. Specifically, I start by looking at the association of income, sales amount characteristics, race, and loan type variables along the chain. Not only can these vacancy chain results be compared to the previous One-Link Analysis, the existence of discrete links that start at new construction offers an opportunity to see how these measures might differ across links: for example, is the Link 0 and Link 1 income relationship (between the people who buy houses from people who bought new houses) different than the Link 2 and Link 3 income relationship? Having a measure of these elasticities, rank-rank slopes, and transition matrices allows one to get a sense of the distributional impact of building an additional new house.

3.3.9.1 Interlink Elasticities

I extend the One-Link Analysis regression to cover different links. Specifically, it is the slope coefficient β from the following regression:

$$\log(Y_{c,l}) = \alpha + \beta \log(Y_{c,l-1}) + \epsilon_{c,l}$$

where c represents a distinct vacancy chain, and l represents the link in the chain. Hence, by our definition of chains, the $l - 1$ link represents the characteristics of the seller who sold a house to buyer l. As a reminder, a value of β close to 1 represents a perfect correlation of a particular characteristic, and a value of β close to 0 represents no correlation (e.g., completely random house search behavior).

Table 3.12: Vacancy Chain Elasticity Estimates: Income and Sales Amount

Dependent Variable: Characteristics of Buyer

	(1) Log Median Tract Income	(2) Log Median Tract Income	(3) Log Household Income	(4) Sale Amount	(5) Sale Amount
Log Median Tract Income (Seller)	0.246	0.248			
	(0.012)	(0.027)			
Log Household Income (Seller)			0.303		
			(0.028)		
Sale Amount (Seller)				0.353	0.448
				(0.015)	(0.029)
Constant	8.364	8.411	8.017	8.357	7.131
	(0.132)	(0.305)	(0.325)	(0.197)	(0.380)
Observations	6786	1162	1162	6786	1162
Sample	Full Chain	Linked HMDA	Linked HMDA	Full Chain	Linked HMDA

Standard errors in parentheses

The results in Table 3.12 show that, when measuring the median tract income of buyers and sellers, the estimated interlink elasticity is about 0.25. This estimate does not change very much when moving between the sample with the full chain or only the linked HMDA sample. Using the household specific income in the matched HMDA sample only, the estimated elasticity is 0.3, a higher elasticity than estimated using median tract income. This is similar to the results for the One-Link analysis. Again, it suggests that using average tract characteristics instead of household specific characteristics to measure interlink elasticities adds an important nuance.

To put these estimates into more context, I again use the thought experiment similar to the one described in Deutscher and Mazumder (2021) to calculate how many links needed for an outlier to regress near to the mean of the sample in income, given the estimated elasticities. A 0.25 elasticity of income implies that a house bought by a household who has a log difference of 1 (about 2.71x average income) should take, on average, about 3 links to reach a household within 0.02 log difference of average sample income. Similarly, if the elasticity is 0.3, then it would take 4 links on average. Hence, even though there is a possible downward bias of the elasticity coefficient using median tract income data instead of household income data, the results are qualitatively similar in terms of the economic implications for interlink mobility.

Similar to the One-Link Analysis, the estimated elasticities for sales amount look different. A 0.35 or 0.45 elasticity of sales amount implies that a house bought by a household who has a log difference of 1 above average (about 2.71x average income) should take, on average, about 4-5 links to reach a household within 0.02 log difference of average sample sales amount. I interpret the reasons for the higher elasticity estimates compared to income to be no different than the reasons stated for the results associated with Table 3.3.

Table 3.13: Vacancy Chain Analysis: Income Elasticity by Link

Dependent Variable: Characteristics of Buyer

	(1) Log Median Tract Income	(2) Log Median Tract Income	(3) Log Household Income
Log Median Tract Income (Seller)	0.206	0.264	
	(0.020)	(0.047)	
Log Median Tract Income (Seller) * Link=2	0.060	0.015	
	(0.029)	(0.066)	
Log Median Tract Income (Seller) * Link=3	0.050	-0.063	
	(0.029)	(0.067)	
Log Household Income (Seller)			0.290
			(0.045)
Log Household Income (Seller) * Link=2			0.010
			(0.069)
Log Household Income (Seller) * Link=3			0.031
			(0.066)
Observations	6786	1162	1162
Sample	Full Chain	Linked HMDA	Linked HMDA

Standard errors in parentheses

Table 3.13 shows the results for interlink income elasticities, but the coefficients are interacted by link. Hence, the excluded category represents the elasticities estimated going from link 0 to link

1 (buyers who bought the house that was sold by the seller who bought into new construction), and the other interactions are the estimated elasticities relative to the first link. The take-away is that there does not seem to be a statistically significant difference across links, except that the second and third link's elasticities might be higher than the first link's elasticities when using tract income measures with the Full Chain sample. This may be due to real differences in interlink elasticities as buyers/sellers go along the links, perhaps due to new construction attracting a more diverse demographic or thicker market. However, there is no direct evidence of this.

To explore the robustness of these results further, I investigate the idea that changes in the distribution (i.e., variance) of the distribution of link characteristics might account for differences of these estimates across links. Deutscher and Mazumder (2021) write that rank-rank slope estimates are closely related to elasticity estimates, but are more robust to variance or other distributional changes. The results in rank-rank slope are presented in Appendix C.3. The qualitative results do not change because the rank-rank slope estimates are remarkably similar to the elasticity estimates, but there are technically no longer statistically significant differences between the rank-rank slope estimates across the links compared to Link 0.

3.3.9.2 Vacancy Chain Analysis: Transition Matrices

Like in the previous sections, I use transition matrices to characterize how racial group and loan-type characteristics are correlated across the links of the vacancy chains. Since these vacancy chains have multiple links, there are two ways of calculating transition matrices. The first is to use all of the links and calculate an overall transition matrix; if the process is Markov, then transition rates should not vary based on which link it is on the chain. This overall transition matrix will be labeled the multi-link transition matrix, since it is an aggregation of multiple links in the chains. The alternative is to calculate separate transition matrices from each link (link 0 to link 1, link 1 to link 2, etc.); this allows for any idiosyncratic differences of newer housing to be captured, but possibly at the cost of lower sample size (and thus higher small sample variance in the estimation of these matrices). This analysis will be labeled link-specific transition matrices. The overall transition matrix is shown in Table 3.14. The link-specific transition matrices are reported in Appendix C.5.

Similar to the One-Link Analysis, the multi-link transition matrix has significant diagonal elements but also significant off-diagonal elements. Although the off-diagonal elements in this analysis may be slightly smaller than those in the corresponding transition matrices in the One-Link Analysis, the results for reaching steady state are qualitatively similar. By Link 3, it is very close to steady state.

I report the impulse-response analysis where I assume new housing is built in the racial or loan type category of the predominant group and see how long it takes for such housing to reach

Table 3.14: Vacancy Chain Transition Matrix: Race

Buyer Race \ Seller Race	Non-Hispanic White	Hispanic White	Black	Other Minority
Non-Hispanic White	0.5076	0.2756	0.2432	0.2829
Hispanic White	0.2203	0.5028	0.4595	0.2465
Black	0.0173	0.0341	0.0811	0.0308
Other Minority	0.2548	0.1875	0.2162	0.4398

Table 3.15: Vacancy Chain Transition Matrix: Loan Type

Buyer Loan Type \ Seller Loan Type	*Conv*	*FHA/VA*
Conv	0.73	0.53
FHA/VA	0.27	0.47

different parts of the population. The analysis is done with both the overall multi-link transition matrix and the link-specific transition matrices. Although there are quantitative differences, the overall pattern of quickly returning close to steady state is very similar to the One-Link Analysis. However, the implied steady state population of these matrices differ from the One-Link analysis because this is using a much more selected sample (due to matching across the different links of the chain and due to conditioning on new housing for Link 0). Interestingly, this selection seems to suggest that the steady states implied by these matrices have more Hispanics than the overall population, suggesting faster diffusion across that ethnic category.

Comparing the multi-link matrix analysis to the link-specific matrices analysis, there are differences in the patterns of these responses across the links of the chain. However, by Link 3, the resulting distributions from the two methods are not too different from each other.

Table 3.16: Vacancy Chains Analysis (multi-link transition matrix): Race Impulse Response

Link	0	1	2	3	...	∞	Actual
Non-Hispanic White	100%	51%	39%	37%		36%	46%
Hispanic White	0%	22%	29%	32%		33%	23%
Black	0%	2%	3%	3%		3%	3%
Other Minority	0%	25%	29%	29%		28%	28%

Table 3.17: Vacancy Chains Analysis (link-specific transition matrices): Race Impulse Response

Link	0	1	2	3	...	∞	Actual
Non-Hispanic White	100%	48%	37%	39%		-	46%
Hispanic White	0%	28%	31%	27%		-	23%
Black	0%	2%	1%	4%		-	3%
Other Minority	0%	23%	31%	30%		-	28%

Table 3.18: Vacancy Chains Analysis (multi-link transition matrix: Loan Type Impulse Response

Link	0	1	2	3	...	∞	Actual
Conventional	100%	73%	68%	67%		66%	64%
FHA/VA	0%	27%	32%	33%		34%	36%

Table 3.19: Vacancy Chains Analysis (link-specific transition matrices): Loan Type Impulse Response

Link	0	1	2	3	...	∞	Actual
Conventional	100%	74%	70%	64%		-	64%
FHA/VA	0%	26%	30%	36%		-	36%

3.3.10 Robustness of Results to Matching Strictness

One of the biggest contingencies for the accuracy of these results is the quality of the matching algorithm. To alleviate concerns, I redo the two main analyses of the interlink income elasticities using a variety of Jaro-Winkler distances (described in Winkler (1990)); this varies the key measure of how strict the matching algorithm is to fuzzy matching on the first name. The results for the default Jaro-Winkler distance of 0.03 is reported again in the middle for easy comparison.

Table 3.20: One-Link Analysis: Interlink Elasticities with Alternative Jaro-Winkler Distances

Jaro-Winkler Dist	Interlink Tract Income Elasticity (SE)	Interlink HH Income Elasticity (SE)	N (full)	N (HMDA-matched)
0.01	0.288	0.365	60713	15927
	(0.004)	(0.007)		
0.02	0.294	0.370	63867	17050
	(0.004)	(0.007)		
0.03	0.296	0.364	69979	18891
	(0.004)	(0.007)		
0.05	0.299	0.363	77302	21097
	(0.003)	(0.007)		
0.10	0.280	0.350	94291	25964
	(0.003)	(0.006)		
0.15	0.269	0.341	100176	27458
	(0.003)	(0.006)		

Note: 0.03 is the value used in the baseline results.

Quite remarkably, the estimates for the baseline results are relatively stable for any Jaro-Winkler distance less than 0.05. I use a distance of 0.03 for conservatism in terms of being relatively sure of accurate matches. However, as these distances increase, the number of matches increases but the estimated elasticity goes closer to zero. This is all consistent with poor quality matching adding noise to these variables. Overall, these alternative Jaro-Winkler distances suggest that the results are not too influenced by arbitary cutoffs in terms of the parameters in the matching algorithm.

3.4 Conclusion

Vacancy chains reveal a lot of information about the nature of churning in housing markets and how vacancies that arise due to new inventory can have cascading effects. This paper tries to characterize those effects, studying whether vacancy chains that might start with a select socioeconomic group of homeowners could have spillover effects on other homeowners through these vacancy chains. To the extent that the same socioeconomic groups (whether characterized by income, neighborhood, house value, or race) sell to the same socioeconomic groups, such spillover results may be limited. But this paper reinforces the general conclusions of the existing literature on housing chains with several nuances: that is, even though there is some persistence of socioeconomic characteristics through the different links in the vacancy chains, these correlations are not strong enough to prevent regression to the mean three to six links down the chain. For most of the results, three or four links down the chain is enough to regress to the mean (in expectation). The added nuance is that using household level characteristics, rather than tract-level characteristics, does tend to result in higher estimates of persistence (or interlink elasticity). In some borderline cases, for a chain that starts with very expensive housing, this may add one more link to the amount needed to regress sufficiently close to the mean. This nuance does not change the overall conclusion that building new housing can have wide-ranging effects on vacancies.

APPENDIX A

Chapter 1 Appendix

A.1 Demand for House Size by Family Size

This section establishes basic background facts about the empirical relationships between house size, family size, income, and age. Demand for house size is measured in the Census as bedrooms. To understand the correlation between house size and family size, I regress house size (number of bedrooms) on the reported household size while controlling for a variety of age and income variables, as well as location fixed effects.

$$N_{ij} = \alpha_j + \beta HHsize_{ij} + \lambda X_{ij} + \epsilon_{ij}$$

where N is number of bedrooms, HH_{size} is household size, and X contains age, age squared, income, income squared. There are state-urban pair fixed effects.

Table A.1: Bedrooms and Household Size

Dependent Variable: Number of Bedrooms

	(1) 1960 Census	(2) 1970 Census	(3) 1980 Census	(4) 1990 Census	(5) 2000 Census	(6) 2010 Census
Household Size	0.215	0.235	0.166	0.189	0.182	0.221
	(0.010)	(0.010)	(0.008)	(0.010)	(0.010)	(0.006)
Observations	490759	550583	3625320	860730	4733176	1087148
State Urban Fixed Effects	x	x	x	x	x	x
X Covariates	x	x	x	x	x	x

The regression is estimated separately for each cross-section in the Census from 1960 to 2010. The results show that there is a strong positive relationship between household size and house size, and this relationship is consistent over time. This result is in line with the theoretical framework in papers like Banks et al. (2015), where larger households have larger housing needs and therefore higher demand for larger housing. It is with this framework I analyze housing demand over time and why houses are getting bigger even though American families are getting smaller.

A.2 The Historical Puzzle

This section establishes that observable household characteristics, particularly growth in real income, cannot adequately explain the total historical growth in housing size, I turn to the Census data, which asks households to report the number of bedrooms in their house. I assume bedrooms are good proxies for a more direct measure like square footage. However, from the American Housing Survey, there is reason to suppose that square footage of per room has been increasing. As such, the following analysis will understate the extent to which real income cannot fully explain growth in housing size.

A.2.1 Engel Curve Estimation

The estimating equation for the relationship between income and housing demand in 1960 follows the standard Engel elasticity estimation. If the relationship between log quantity (bedrooms) and log income is greater than unity, then the good is a luxury good. Conversely, if the elasticity is less than unity, then the good is a necessity good. The baseline estimating equation is below:

$$\log(N_{ic}) = \alpha_c + \beta \log(Y_{ic}) + \lambda X_{ic} + \epsilon_{ic}$$

where α_c are location c fixed effects, N is bedrooms, and Y is real income. Locations are state-urban/rural pairs, with urban areas further subdivided into areas considered inside a principal city, outside the principal city, or mixed. In essence, these regressions should capture the existing local area relationship between income and housing size demand in 1960, including curvature features of the income-demand curve, like luxury or necessity good features, conditional on the metropolitan status of the area and regional variation as measured on the state level.

The estimated elasticities are consistently well below one. With and without fixed effects, they are very similar, which suggests that regional and urban/rural differences are not driving the curvature. As demographic variables are included, the curvature of the Engel curve is even more apparent. This is strong evidence that house size, as measured by the number of bedrooms, is a necessity good. That is, at higher levels of income, the growth in housing size slows down, which may reflect the fact that there are luxury substitutes for housing size (like granite countertops or better locations).

Table A.2: Engel Elasticities

Dependent Variable: Log Number of Bedrooms

	(1)	(2)	(3)	(4)
Log(Income)	0.102	0.110	0.072	0.077
	(0.001)	(0.001)	(0.001)	(0.001)
Observations	482749	482749	480432	480432
Location FE		x		x
Demographic Controls			x	x

Standard errors in parentheses

Using these regression coefficients based on relationships from 1960, I project, for Census years following 1960, estimated log housing size demand for each household based on their given measured characteristics. I then plot the average for each Census year. If the demand relationship between income and housing from 1960 stayed constant, and if changes in the distribution of houses across locations (states + metropolitan status), demographic variables, and income could fully explain the increased demand for larger housing, then we would expect that actual time series and our predicted time series to be similar.

However, the following figures both show that, across a variety of specifications, household observables fall far short of explaining the historical trends. Even the specification with income only, which represents an unconditional Engel curve, shows that income by itself only explains about half of the log point rise in bedroom demand. The inclusion of demographic variables suggests demand was predicted to be flat from the 1960's onwards. The inclusion of interaction terms between income and demographic variables does not seem to change the qualitative result. Neither does estimating location-varying coefficients.

As robustness checks, I use alternative estimates of income from industry and education information, which may better capture the housing decisions based on permanent income. I also check for the influence of censoring on the data since bedrooms larger than 4 are coded as four in the Census. The results are given in Table A.3. Neither of these alternative specifications give different qualitative results.

These empirical results show one main point: the usual explanations of income are inadequate in explaining the rise of house size. In fact, the demographic variables (age and especially household size) are pushing house size down over time. Hence, a combination of preferences (or technology) and price changes must be occurring over time.

Figure A.1: Average Log Number of Bedrooms 1960-2017, Actual vs. Predicted from 1960 Income and Demographic Coefficients

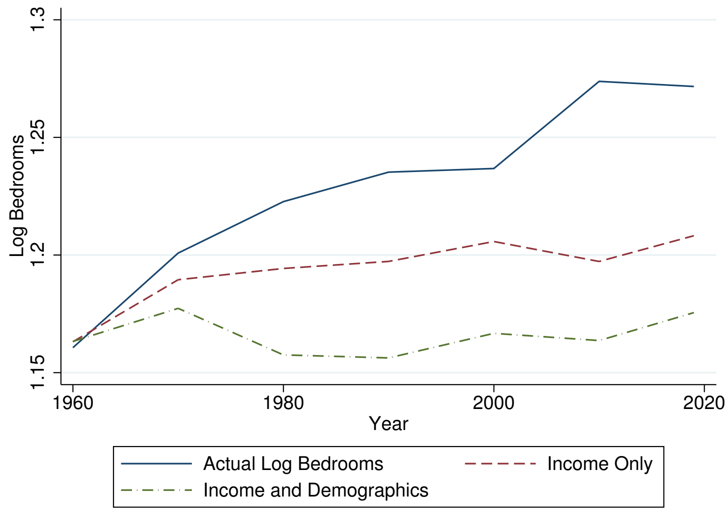

Data from Steven Ruggles, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. IPUMS USA: Version 9.0. Minneapolis, MN: IPUMS, 2019. https://doi.org/10.18128/D010.V9.0

Figure A.2: Average Log Number of Bedrooms 1960-2017, Actual vs. Predicted from 1960 Income and Demographic Coefficients (with Interaction and Location-Vary Terms)

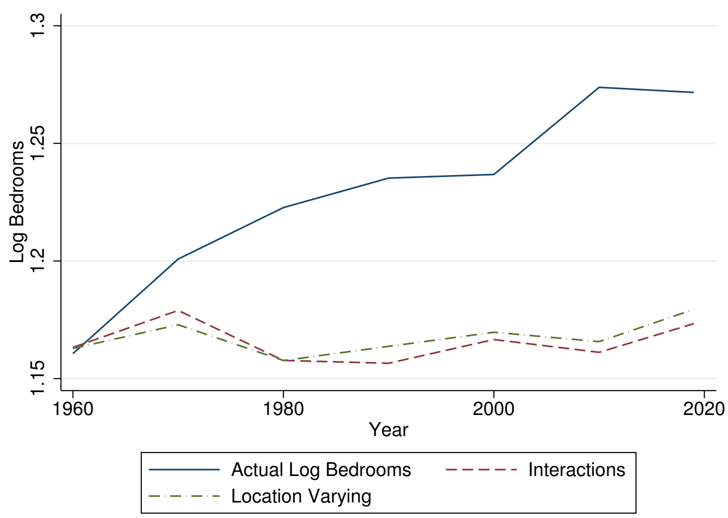

Data from Steven Ruggles, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek. IPUMS USA: Version 9.0. Minneapolis, MN: IPUMS, 2019. https://doi.org/10.18128/D010.V9.0

Table A.3: Engel Elasticities Alternative Estimates

Dependent Variable: Log Number of Bedrooms

	(1) Baseline OLS	(2) Tobit	(3) IV
Log(Income)	0.072	0.076	0.097
	(0.001)	(0.001)	(0.001)
Observations	480432	480432	480432
Demographic Controls	x	x	x

Standard errors in parentheses

A.3 Model Mathematical Details

The optimization problem for household i within a location is the following:

$$\max_{\{c_{it}\},\{h_{it}\}} \sum_{t=0}^{N} \beta^t \left(\frac{c_{it}^{1-\frac{1}{\eta_c}}}{1-\frac{1}{\eta_c}} + \frac{(\theta_{it}\xi_i h_{it})^{1-\frac{1}{\eta_h}}}{1-\frac{1}{\eta_h}} \right) + \zeta_i^L$$

$$s.t. \sum_{t=0}^{N} \frac{c_{it} + p(h_{it})}{(1+r)^t} = M_i$$

where the pricing function is (apparently linearly) given as $p(h_{it}) = p_0 + p_h h_{it}$

I suppress the i subscript for exposition purposes. Let Ω be the Lagrange multiplier on the budget constraint. First order conditions give:

$$(\beta(1+r))^t c_t^{\frac{-1}{\eta_c}} = \Omega \tag{A.1}$$

$$(\beta(1+r))^t (\theta_t \xi)^{1-\frac{1}{\eta_h}} h_t^{\frac{-1}{\eta_h}} = p_h \Omega \tag{A.2}$$

A.3.1 Relative Demand Equations and Numerical Solutions

Cancelling out the Ω terms gives the (Euler) intertemporal optimality conditions and the intratemporal optimality conditions below, respectively.

$$c_t^{\frac{1}{\eta_c}} = \beta(1+r) c_{t-1}^{\frac{1}{\eta_c}} \tag{A.3}$$

$$c_t^{\frac{1}{\eta_c}} = (\theta_t \xi)^{\frac{1}{\eta_h}-1} h_t^{\frac{1}{\eta_h}} p_h \tag{A.4}$$

The two equations above, combined with the budget constraint, are enough to solve the entire household problem. Note that pricing function is consistent with the overall nonlinear structure of the minimum lot size regulation. Hence, the actual solution algorithm loops over multiple prices to see where the solution lies on the budget line. Even without nonlinear pricing, no known analytical solutions exist for these kinds of preferences. The solutions are calculated using a shooting algorithm which repeatedly guesses initial consumption (for the initial period), and then creates the stream of consumption and housing demands consistent with the optimality conditions above. The expenditure of that guess is then compared to the budget constraint. Relevant adjustments are made to the initial guess for consumption based on whether total expenditure is below or above the budget. The solution is found when the budget constraint holds, given some small tolerance.

To understand the relative demand equations, take logs of the intratemporal constraint:

$$\frac{\log h_t}{\eta_h} - \frac{\log c_t}{\eta_c} = -\log p_h + \left(\frac{\eta_h - 1}{\eta_h}\right)\log\theta_t + \left(\frac{\eta_h - 1}{\eta_h}\right)\log\xi \tag{A.5}$$

A.3.1.1 Deriving Engel Curves

Now let us focus on the intratemporal conditions. Let M_t be the optimal expenditure in period t. The budget constraint for that period is:

$$\begin{aligned} M_t &= c_t + p_h h_t && \text{(A.6)} \\ &= \Omega^{-\eta_c}(\beta(1+r))^{\eta_c t} + \Omega^{-\eta_h}(\theta_t\xi)^{\eta_h - 1}(\beta(1+r))^{\eta_h t} && \text{(A.7)} \end{aligned}$$

Writing M_t and Ω in logs gives:

$$\log(M_t) = \log\left[e^{-\eta_c \log\Omega}(\beta(1+r))^{\eta_c t} + e^{-\eta_h \log\Omega}(\theta_t\xi)^{\eta_h - 1}(\beta(1+r))^{\eta_h t}\right] \tag{A.8}$$

Using the implicit function theorem on the equation above implies:

$$\begin{aligned} \frac{d\log\Omega}{d\log M_t} &= \frac{-1}{-\left(\frac{-\eta_c\Omega^{-\eta_c}(\beta(1+r))^{\eta_c t} - \eta_h\Omega^{-\eta_h}(\theta_t\xi)^{\eta_h-1}(\beta(1+r))^{\eta_h t}}{M_t}\right)} && \text{(A.9)} \\ &= \frac{-1}{-\left(\frac{-\eta_c c_t - \eta_h p_h h_t}{M_t}\right)} && \text{(A.10)} \\ &= \frac{-1}{\eta_c s_c + \eta_h s_h} && \text{(A.11)} \end{aligned}$$

where s_c and s_h are the shares of expenditures of each good, respectively, within period t.

Define $\bar{\eta} = \eta_c s_c + \eta_h s_h$, the expenditure weight shares of the parameters η. Finally, going back to the original first order conditions:

$$\frac{d\log c_t}{d\log M_t} = \frac{d\log c_t}{d\log \Omega}\frac{d\log \Omega}{d\log M_t} \tag{A.12}$$

$$= \frac{\eta_c}{\eta_c s_c + \eta_h s_h} \tag{A.13}$$

$$= \frac{\eta_c}{\bar{\eta}} \tag{A.14}$$

An analogous derivation gives:

$$\frac{d\log h_t}{d\log M_t} = \frac{\eta_h}{\bar{\eta}} \tag{A.15}$$

Thus, the parameters η_h and η_c, in relation to each other and at the optimum, govern the shape of the Engel curve.

A.4 Synthetic Control Details

A.4.1 Predictor Variable Weights

Variable weights are estimated by a nested optimization problem defined in the main paper. The relevant weights are given below:

Table A.4: Predictor Variables Weights

Variable	Weights
Minority Population	0.0040
Median HH Income	0.0017
Median Rent	0.0012
Log Square Feet (1991)	0.2400
Log Square Feet (1993)	0.1368
Log Square Feet (1995)	0.2340
Log Square Feet (1997)	0.2933
MSA Pop Growth (1991-1997)	0.0873
Density (1990)	0.0015

Note: Weights are a function of importance and the magnititude of the underlying variables. Minority Population, Median Income, and Median Rent characteristics are tract-level characteristics weighted by housing units built in the pre-period.

A.4.2 Alternative Specifications: Synthetic Control Cities, Weights, and Results

The following shows that the exclusion of certain variables do not significantly change the results. Specifically, I do two alternative versions of the synthetic control method: the first simply drops the MSA population growth and city density values to see if the results change. The idea is to understand how these population variables may be driving the underlying results.

Table A.5: Alternative Specification 1: City Weights

City	Weight
Plano	0.429
San Antonio	0.321
Austin	0.245
Round Rock	0.005

Figure A.3: Alternative Specification 1 for Houston: Synthetic Control, Minimum Lot Size Reduction in 1999

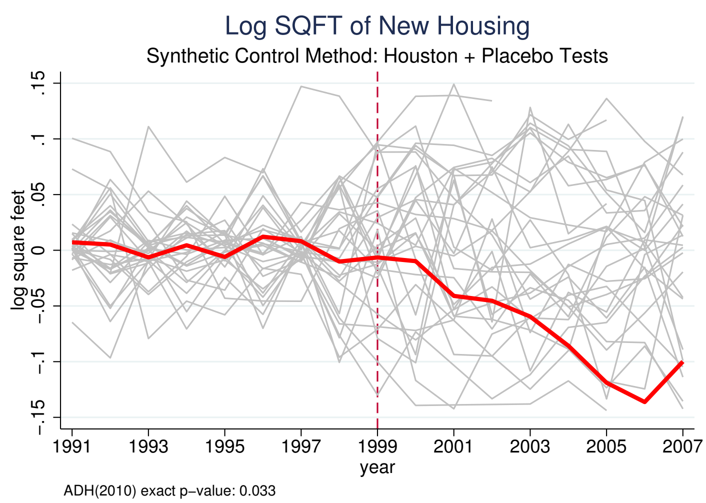

Notes: The synthetic control method chooses a convex combination of control cities. Gray lines are placebo tests where the same synthetic control procedure is repeated for all cities in the donor pool.

The second alternative method is to only use each year's outcome variable (log average square feet) as predictors in the pre-period. This is a simpler approach which imposes the strongest parallel trends and level matching assumptions for the pre-period outcomes.

Lastly, the calculated exact p-value does not change because Houston's ratio of its post-period deviation of house from that of its own synthetically created control city (compared to the pre-period) is still the largest out of all possible placebo cities, strongly suggesting there is a real

Table A.6: Alternative Specification 2: City Weights

City	Weight
Plano	0.460
San Antonio	0.234
Austin	0.222
Fort Worth	0.084

Figure A.4: Alternative Specification 2 for Houston: Synthetic Control, Minimum Lot Size Reduction in 1999

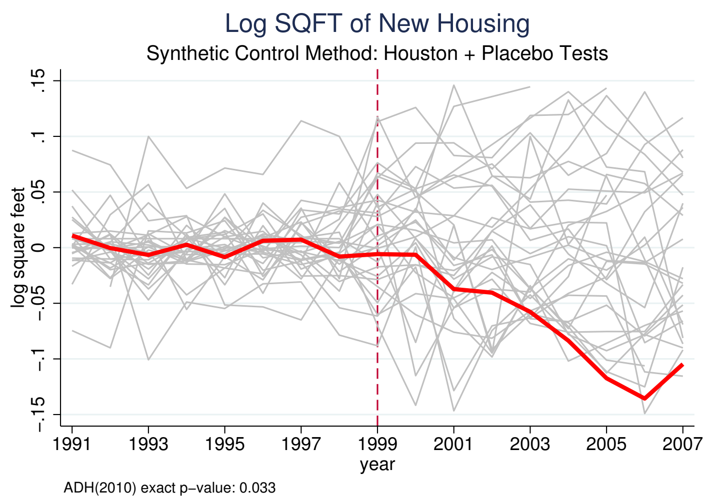

Notes: The synthetic control method chooses a convex combination of control cities. Gray lines are placebo tests where the same synthetic control procedure is repeated for all cities in the donor pool.

structural shift happening in Houston.

A.4.3 Alternative Houston Results: Standard Diff-in-Diff

As a robustness check, I use the standard difference-in-difference empirical setting using all the donor cities. The general empirical framework is given below.

$$Y_{ijt} = \phi_j + \lambda_t + \sum_{k \neq EVENT} \mathbb{1}\{t = k\}\beta_k Houston_j + \epsilon_{ijt}$$

where Y is the outcome variable (log square feet), ϕ_j and λ_t are jurisdiction and year fixed effects, and where $Houston_j$ is an indicator variable for a jurisdiction (Houston) changing their policy. Hence, in the standard event study framework, the coefficients β_t represent the differential level of the outcome variable for jurisdictions (i.e., Houston) that decreased regulations, relative to the jurisdictions that kept their regulations the same.

The standard diff-in-diff results for both the stock and flow of average housing size built each year is plotted below. The relative log square feet of housing built in Houston exhibits parallel trends from 1993 to 1999, but significantly decrease afterwards. The overall change from 1998 to 2006/2007 is about 14 log points, which is even larger than the effect identified using synthetic control methods. However, one may argue there is a pre-trend before 1993, so the baseline synthetic control specification in the paper places the estimate into more context.

Figure A.5: Houston: Diff-in-diff, Minimum Lot Size Reduction in 1999

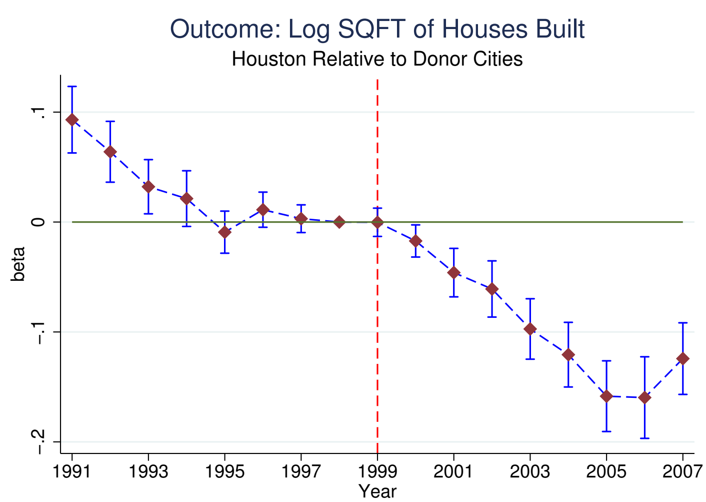

A.4.4 Alternative Simulation: Welfare Effects

The household gains (for households who were always in Houston) in the alternative scenario where the price changes are interpreted as transitory, are less substantial. The shape of the distribution looks very similar because these two scenarios change only in the price inputs, not in the preference draws or underlying distribution of income and household size. But both the magnitudes of the gains (less than $8,000 on average) and the range of the heterogeneity (less than $6,000) are smaller in scale. However, even though these amounts might be less economically relevant, they fully demonstrate that the mechanisms in the model work across different assumptions in the interpretation of the Houston reduced form estimates. The alternative simulations with asset price effects also give similar qualitative conclusions as the baseline but with the average welfare effects and heterogeneity also being a magnitude smaller.

Figure A.6: **Alternative Simulation:** Household Lifetime Gains Across Income and Household Size (2010 Dollars)

(a) Income (b) Household Size

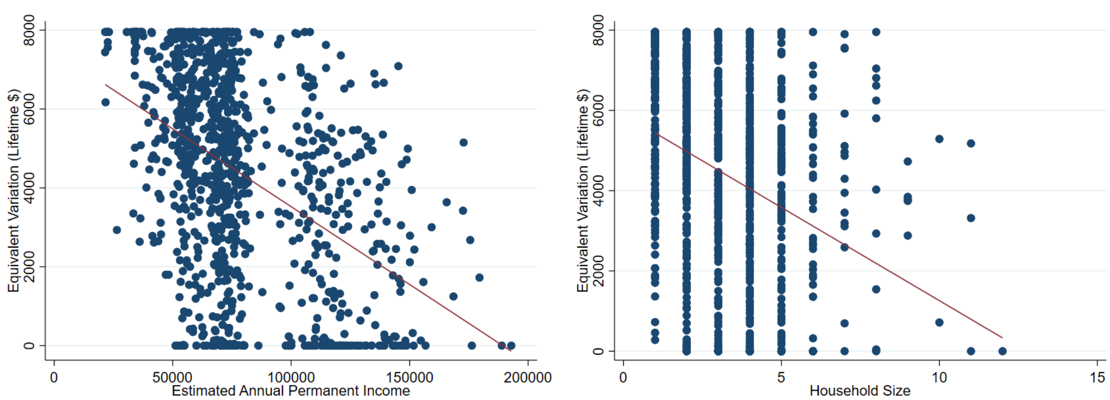

Figure A.7: **Alternative Simulation With Asset Price Effects:** Household Lifetime Gains Across Income and Household Size (2010 Dollars)

(a) Income (b) Household Size

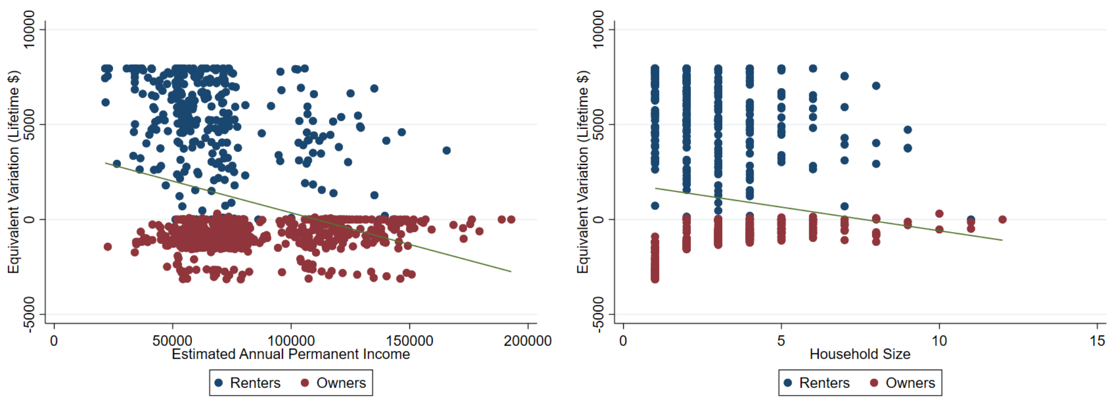

A.5 Robustness of Model to Choice of Calibrated Parameters

This section analyzes how the fundamental hetereogeneity results compare across alternative specifications for the chosen interest rate, the discount rate, and weight parameter (of current income vs peer income in one's education/industry group) used in the calculation of a household's permanent income. Intuitively, the discount rate affects intertemporal consumption smoothing, the interest rate affects consumption smoothing and wealth, and the weight parameter only affects wealth.

In the subsequent series of figures is plotted the analogous baseline results for the equivalent variation for each simulated household drawn from the empirical distribution. For different calibrated parameters used in the model, I perturb them and re-estimate the model completely, running the simulations yet again.[1] The qualitative conclusion is that the average welfare gain from the deregulation event is unstable and can vary thousands of dollars across different specifications. The largest discrepancy between the average welfare gains comes from perturbing the interest rate, which seems to have significant wealth and discounting effects. However, the downward sloping nature of the heterogeneity is persistent; across different specifications, the range of variation (between the highest and lowest welfare gains) is about \$20000 to \$27000. The largest discrepancy for the heterogeneity results across income comes from perturbing the permanent income weighting term w; this means that heterogeneity in welfare gains could be somewhat affected by how permanent income is measured. This result is consistent with the general intuition that weighing idiosyncratic annual income too much (i.e., $w \gg 0$) creates a noisy measure of true permanent

[1]The idiosyncratic shocks are saved (or seeded, in computer science terms) and therefore identical across each run, to offer a comparison that is not affected by small sample issues. As a result, the distribution of results in the subsequent figures looks very similar across different runs.

income; hence, the resulting slope of welfare may suffer from an attenuation bias. Finally, the results do not seem to be qualitatively affected by the discount rate, suggesting that the results are not driven by assumptions about households' impatience.

Figure A.8: Robustness: Welfare Heterogeneity Results Across Income, For Different Calibrated Parameters

(a) $\beta = 0.92$

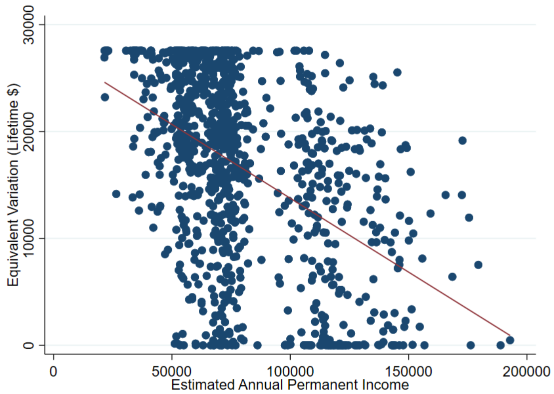

(b) $\beta = 0.88$

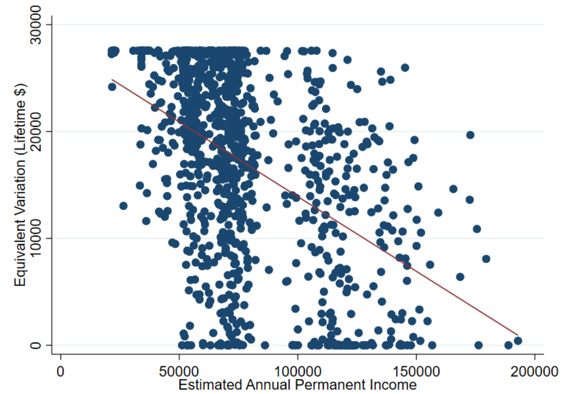

(c) $r = 0.12$

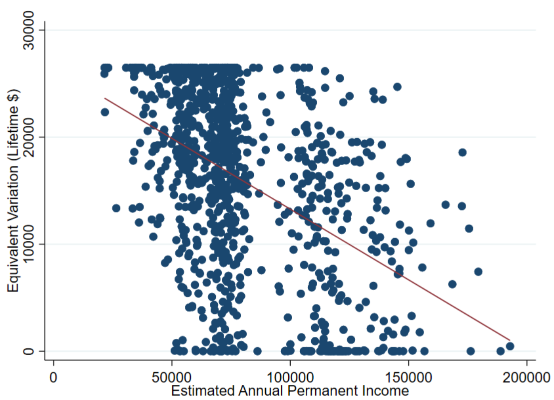

(d) $r = 0.08$

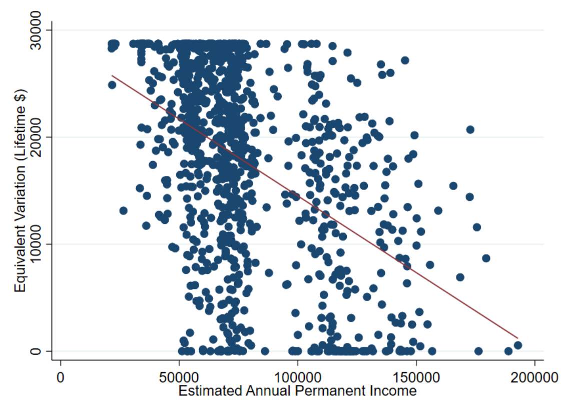

(e) $w = 0.05$

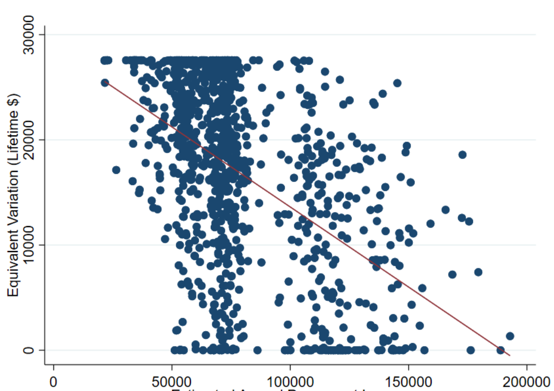

(f) $w = 0.25$

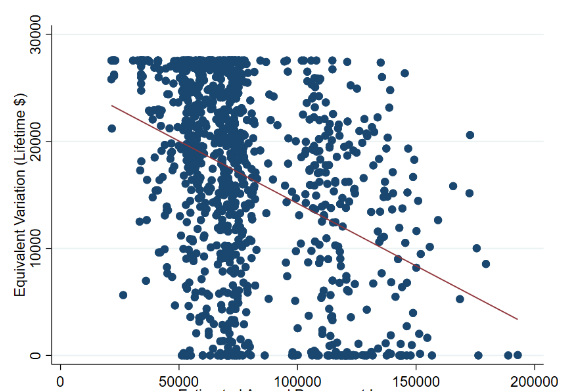

Figure A.9: Robustness: Welfare Heterogeneity Results Across Family Size, For Different Calibrated Parameters

(a) $\beta = 0.92$

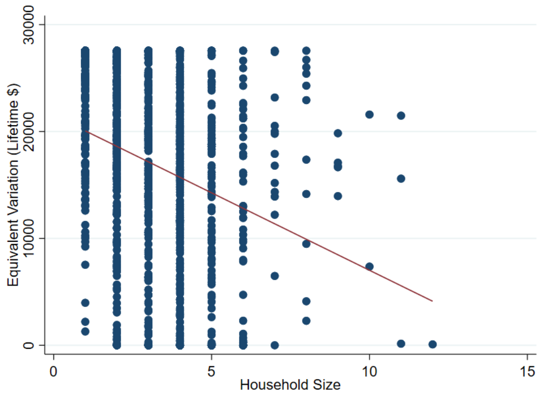

(b) $\beta = 0.88$

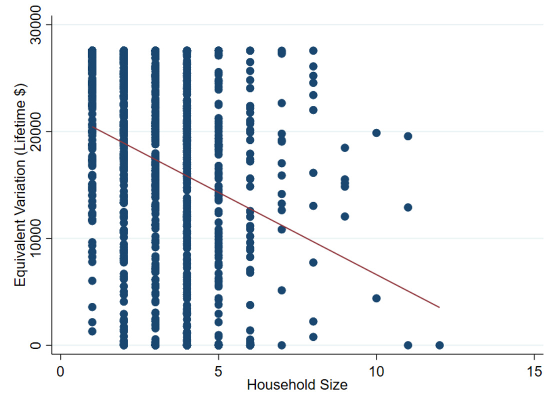

(c) $r = 0.12$

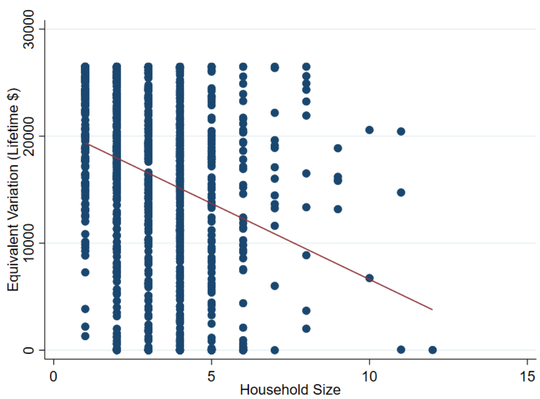

(d) $r = 0.08$

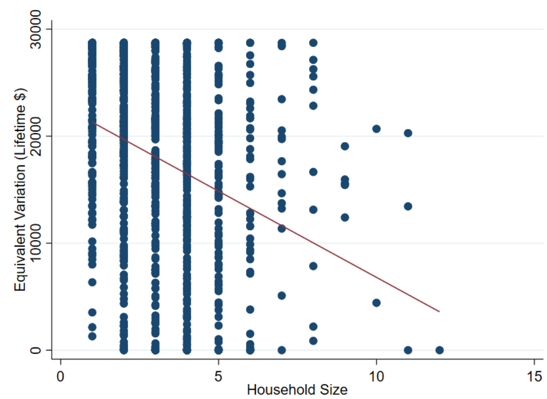

(e) $w = 0.05$

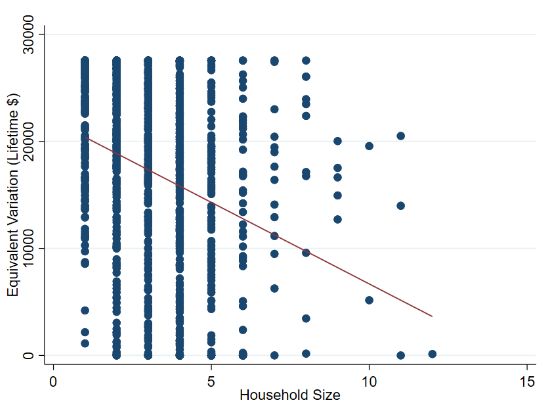

(f) $w = 0.25$

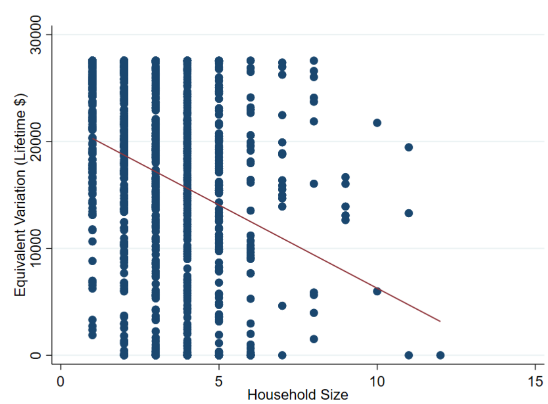

A.6 Roy Model Sorting Mechanisms

A.6.1 The Canonical Roy Model

The Canonical Roy Model [2] models wages and assumes joint normality of the underlying deviation from the wage means in each respective location. Under these assumptions, a standard theoretical result is that if the correlation between the two deviation terms are sufficiently high, what governs negative or positive selection depends on the relative variance of each deviation term. This statement is formalized below:

Suppose source location 0 gives payoff $w_0 + \epsilon_0$ and destination location 1 gives payoff $w_1 + \epsilon_1$ where ϵ_0 and ϵ_1 are jointly normal with standard deviations σ_0 and σ_1 and correlation ρ. If $\rho >> 0$, then there will be positive selection into the destination location if and only if $\sigma_1 > \sigma_0$. There will be negative selection if and only if $\sigma_1 < \sigma_0$.

A.6.2 Housing Size Roy Model

In terms of the model in this paper, destination location 1 is the city which reduces its minimum lot size. Source location is represented by location 0 which is the status quo city (i.e., the city without the policy change).

As such, lifetime utility for household i in location 0 or 1 is given by:

$$u_{0,i} = U(p^0, \theta_i; \omega) + \epsilon_{0_i}$$
$$u_{1,i} = U(p^1, \theta_i; \omega) + \epsilon_{1_i}$$

where p^L are the vector of prices in each location, θ_i is a vector of characteristics with family size, income, and age, and ϵ_L are independent and identically distributed preference terms for each location. ω is a vector of parameters which will be suppressed for exposition purposes.

A first order approximation around the mean values of θ in the population gives:

$$u_{0,i} = U(p^0, \bar{\theta}) + \frac{\partial U(p^0, \bar{\theta})}{\partial M}(M_i - \bar{M}) + \frac{\partial U(p^0, \bar{\theta})}{\partial H}(H_i - \bar{H}) + \epsilon_{0_i}$$
$$u_{1,i} = U(p^1, \bar{\theta}) + \frac{\partial U(p^1, \bar{\theta})}{\partial M}(M_i - \bar{M}) + \frac{\partial U(p^1, \bar{\theta})}{\partial H}(H_i - \bar{H}) + \epsilon_{1_i}$$

[2] "Lecture Note: Self Selection – The Roy Model." David Autor. MIT. Accessed January 18, 2022. https://economics.mit.edu/files/551

Assuming joint normality of H and M, the variance of each term is given below:

$$\sigma_{0,i}^2 = (\frac{\partial U(p^0,\bar{\theta})}{\partial M})^2\sigma_M^2 + (\frac{\partial U(p^0,\bar{\theta})}{\partial H})^2\sigma_H^2 + 2(\frac{\partial U(p^0,\bar{\theta})}{\partial M})^2\sigma_{HM} + \sigma_\epsilon^2$$
$$\sigma_{1,i}^2 = (\frac{\partial U(p^1,\bar{\theta})}{\partial M})^2\sigma_M^2 + (\frac{\partial U(p^1,\bar{\theta})}{\partial H})^2\sigma_H^2 + 2(\frac{\partial U(p^1,\bar{\theta})}{\partial M})^2\sigma_{HM} + \sigma_\epsilon^2$$

Taking the difference of the two equations:

$$\begin{aligned}\sigma_{1,i}^2 - \sigma_{0,i}^2 = \sigma_M^2\big[(\frac{\partial U(p^1,\bar{\theta})}{\partial M})^2 - (\frac{\partial U(p^0,\bar{\theta})}{\partial M})^2\big] + \sigma_H^2\big[(\frac{\partial U(p^1,\bar{\theta})}{\partial H})^2 - (\frac{\partial U(p^0,\bar{\theta})}{\partial H})^2\big] \\ + 2\sigma_{HM}\big[(\frac{\partial U(p^1,\bar{\theta})}{\partial M})(\frac{\partial U(p^1,\bar{\theta})}{\partial H}) - (\frac{\partial U(p^0,\bar{\theta})}{\partial M})(\frac{\partial U(p^0,\bar{\theta})}{\partial H})\big]\end{aligned}$$

A fundamental assumption consistent with this paper is that decreases in the price of house size decreases the returns to additional income or household size. This is the underlying heterogeneity channel that affects the direction of selection. That is, assume:

$$\frac{\partial U(p^1,\bar{\theta})}{\partial F} < \frac{U(p^0,\bar{\theta})}{\partial F}$$

for any variable F in ω.

Roy Model Proposition: In first order terms, there is negative selection if and only if $\sigma_{HM} > \frac{-\sigma_M^2\left[(\frac{\partial U(p^1,\bar{\theta})}{\partial M})^2-(\frac{\partial U(p^0,\bar{\theta})}{\partial M})^2\right]-\sigma_H^2\left[(\frac{\partial U(p^1,\bar{\theta})}{\partial H})^2-(\frac{\partial U(p^0,\bar{\theta})}{\partial H})^2\right]}{2\left[(\frac{\partial U(p^1,\bar{\theta})}{\partial M})(\frac{\partial U(p^1,\bar{\theta})}{\partial H})-(\frac{\partial U(p^0,\bar{\theta})}{\partial M})(\frac{\partial U(p^0,\bar{\theta})}{\partial H})\right]}$.

Note that the term on the right is negative. The proposition here is that there is negative selection (i.e., the people with lower than average utility move into city 1) when the correlation between household size and income is not too negative. In reality, because children are normal goods, this condition tends to be satisfied. This negative selection translates monotonically into lower income and smaller families moving into the deregulated city.

APPENDIX B

Chapter 2 Appendix

B.1 Optimal Housing Decisions

To start, here are the wealth equations and conjectured rental and housing asset prices:

$$w_e = y_t - \lambda r_t + p_t h_{t-1} - p_{t-1} h_{t-1} \frac{1}{\beta} + r_t h_{t-1} \tag{B.1}$$

$$r_t = \frac{y_t}{\lambda} + \bar{r} \tag{B.2}$$

$$p_t = \frac{\beta}{1-\beta}(\frac{y_t}{\lambda} + \bar{r} - \rho) \tag{B.3}$$

Basic substitution of the rental price r_t and asset price p_t gives:

$$w_e = y_t - \lambda(\frac{y_t}{\lambda} + \bar{r}) + \frac{\beta}{1-\beta}(\frac{y_t}{\lambda} + \bar{r} - \rho)h_{t-1} - \frac{\beta}{1-\beta}(\frac{y_{t-1}}{\lambda} + \bar{r} - \rho)h_{t-1}\frac{1}{\beta} + (\frac{y_t}{\lambda} + \bar{r})h_{t-1} \tag{B.4}$$

$$= -\lambda\bar{r} + \frac{\beta}{1-\beta}(\frac{y_{t-1} + \mu_t}{\lambda} + \bar{r} - \rho)h_{t-1} - \frac{\beta}{1-\beta}(\frac{y_{t-1}}{\lambda} + \bar{r} - \rho)h_{t-1}\frac{1}{\beta} + (\frac{y_{t-1} + \mu_t}{\lambda} + \bar{r})h_{t-1} \tag{B.5}$$

$$= -\lambda\bar{r} + h_{t-1}\left(\frac{\mu_t}{(1-\beta)\lambda} + \rho\right) \tag{B.6}$$

$$\tag{B.7}$$

Taking expectations and variances of the above gives:

$$E[w_e] = -\lambda\bar{r} + \rho h_{t-1} \tag{B.8}$$

$$V[w_e] = \frac{h_{t-1}^2 \sigma^2}{\lambda^2(1-\beta)^2} \tag{B.9}$$

Substituting into the quadratic utility function:

$$E[U] = -(-\lambda\bar{r} + \rho h_{t-1} - \bar{w})^2 - \frac{h_{t-1}^2 \sigma^2}{\lambda^2(1-\beta)^2} \tag{B.10}$$

Taking first order condition of the above expression with respect to h_{t-1} gives the optimal choice of housing ownership:

$$\begin{aligned}
0 &= -2(-\lambda\bar{r} + \rho h_{t-1} - \bar{w})\rho - 2\frac{h_{t-1}\sigma^2}{\lambda^2(1-\beta)^2} \\
&= (-\lambda\bar{r} - \bar{w})\rho + \rho^2 h_{t-1} + \frac{h_{t-1}\sigma^2}{\lambda^2(1-\beta)^2} \\
&= (-\lambda\bar{r} - \bar{w})\rho + h_{t-1}\left(\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}\right)
\end{aligned}$$

$$\implies h_{t-1}^* = \frac{\rho(\lambda\bar{r} + \bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}} \tag{B.11}$$

B.2 Deriving $F(\bar{r})$

Next I derive the equation for $F(\bar{r})$ given in Chapter 2. I start by writing out utility without the idiosyncratic preference term. Then I plug in the optimal values for h_{t-1}^*.

$$\begin{aligned}
-(E[w_e] - \bar{w})^2 - V(w_e)) &= -(-\lambda\bar{r} + \rho h_{t-1}^* - \bar{w})^2 - \frac{(h_{t-1}^*)^2\sigma^2}{\lambda^2(1-\beta)^2} \\
&= -\left(-\lambda\bar{r} + \frac{\rho^2(\lambda\bar{r} + \bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}} - \bar{w}\right)^2 - \frac{\left(\frac{\rho(\lambda\bar{r}+\bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\right)^2 \sigma^2}{\lambda^2(1-\beta)^2} \\
&= -\left(-(\lambda\bar{r} + \bar{w}) + \frac{\rho^2(\lambda\bar{r} + \bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\right)^2 - \left(\frac{\rho(\lambda\bar{r} + \bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\right)^2 \frac{\sigma^2}{\lambda^2(1-\beta)^2} \\
&= -\left(\left(\frac{\rho^2}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}} - 1\right)(\lambda\bar{r} + \bar{w})\right)^2 - \left(\frac{\rho(\lambda\bar{r} + \bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\right)^2 \frac{\sigma^2}{\lambda^2(1-\beta)^2} \\
&= -\left(1 - \frac{\rho^2}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\right)^2 (\lambda\bar{r} + \bar{w})^2 - \left(\frac{\rho(\lambda\bar{r} + \bar{w})}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\right)^2 \frac{\sigma^2}{\lambda^2(1-\beta)^2} \\
&= F(\bar{r})
\end{aligned}$$

B.3 Deriving the Relationship Between N and $\bar{r}$

With the expression above, the equilibrium condition $F(\bar{r}) + Z^{-1}(1-N) = \mu_0$ can be simplified as following:

$$\bar{A}x^2 = \mu_0 - Z^{-1}(1-N) \tag{B.12}$$

where

$$\bar{A} = \left[-(\frac{\rho}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}})^2 \frac{\sigma^2}{\lambda^2(1-\beta)^2} - \left(1 - \frac{\rho^2}{\rho^2 + \frac{\sigma^2}{\lambda^2(1-\beta)^2}}\right)^2 \right] \tag{B.13}$$

and

$$x = \lambda\bar{r} + \bar{w} \tag{B.14}$$

Note that, from above, $\bar{A}$ is negative. This means that a solution exists if and only if μ_0 is sufficiently small (and possibly negative). This makes sense because expected utility is negative by construction, and for people to live in the city, the outside countryside option has to be worse for at least some people. Solving for $\bar{r}$, this implies:

$$\bar{r} = F^{-1}(\mu_0 - Z^{-1}(1-N)) = \frac{\sqrt{Z^{-1}(1-N) - \mu_0}}{\lambda\sqrt{-\bar{A}}} - \frac{\bar{w}}{\lambda} \tag{B.15}$$

Since the inverse cumulative distribution function has to be increasing in N, the relevant negative in the function above implies the entire function is decreasing in N. This is the mathematical result for which a larger city size decreases rents: a larger city means that the marginal person values it less.

Lastly, it is helpful in the simulations to solve for $\frac{d\bar{r}}{dN}$:

$$\frac{d\bar{r}}{dN} = \frac{-1}{2\lambda\sqrt{-\bar{A}}\sqrt{Z^{-1}(1-N) - \mu_0}} \frac{dZ^{-1}(1-N)}{dN} \tag{B.16}$$

B.4 Example Simulations of Model

The following come from simulations after fixing parameter values[1] and varying housing needs or family size parameter λ. Plotted are the marginal benefit and marginal cost tradeoffs that determine the size of the city in a political economy equilibrium. Note that the Marginal Benefit curves can, in general, shift because of different values of λ. However, in this simulation, the shifts are so slight that the curves are practically the same line.

Figure B.1: Sizes of Cities in Political Economy Equilibrium: MB vs MC of Increasing City Size

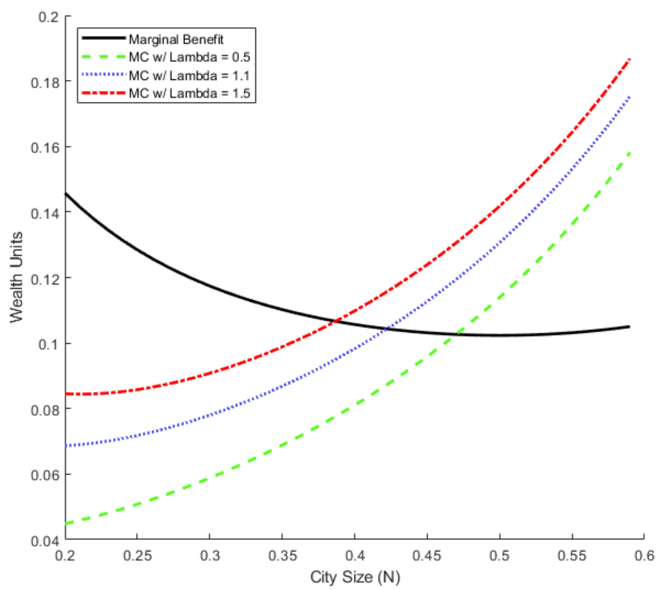

In the simulations above, as housing needs λ increases, the size of the city decreases. With these simulations, the equilibrium size of the city ranges from 0.38 to 0.47.

[1] $\rho = 0.1$, $\sigma = 3$, $\mu_0 = -15000$, $\bar{w} = 100$, $\beta = 0.8$. I assume the idiosyncratic preference term ϵ is distributed normal with a mean of zero and a standard deviation of 10. I assume that $f'(N) = \frac{0.1}{1-x} - 0.1$

B.5 Housing Regulation Distribution Maps

Figure B.2: Distribution of Various Measures of Housing Regulations

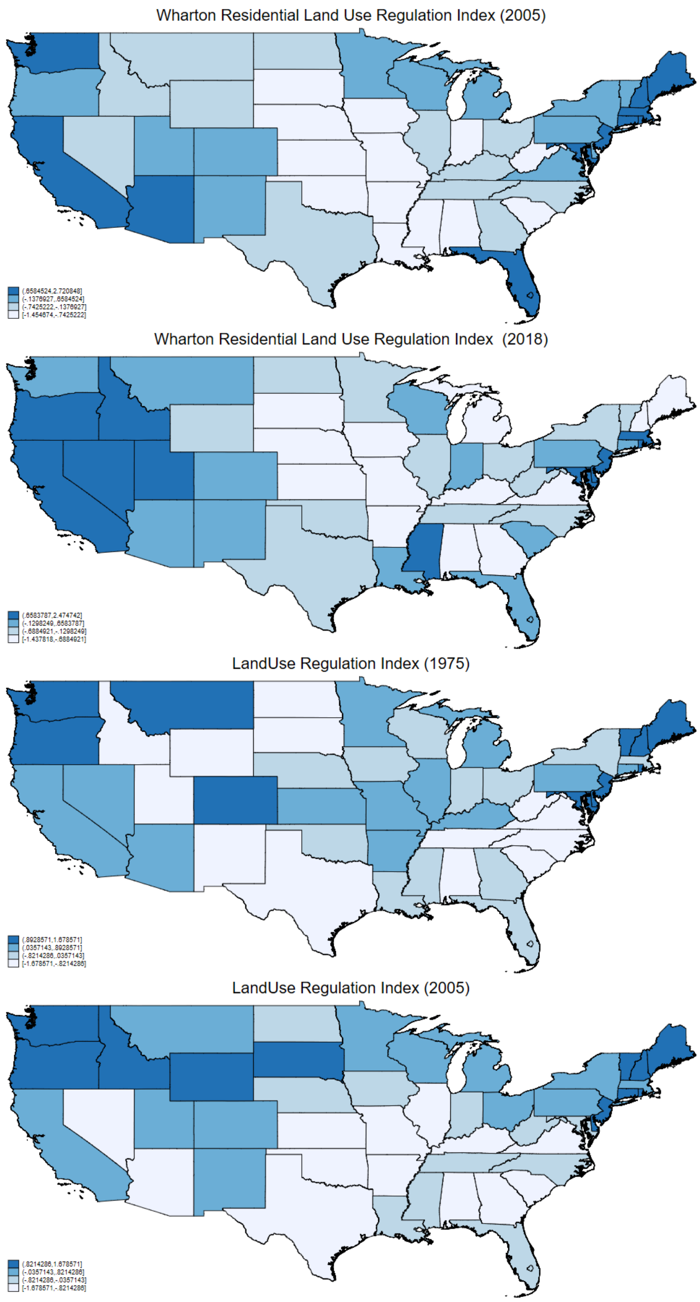

Figure B.3: Distribution of Various Measures of Housing Regulations (continued)

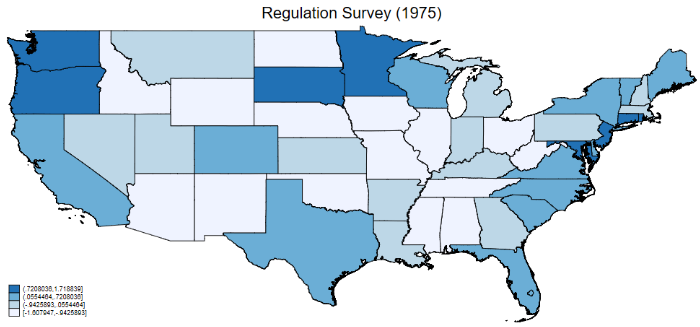

B.6 Robustness: Summary Statistics Results with the Smaller 135 MSA Sample

The following gives summary statistics and main regression results with the smaller 135 MSA sample. The results are largely consistent with the larger sample.

Table B.1: Summary Statistics: MSA Level

	Mean	SD	Min	Max	N
Δ Fertility Rate (1940-1960)	0.64	0.28	-0.92	1.12	135
WWII Casualty Rate (%)	0.24	0.04	0.13	0.52	135
Average Income	850	176	264	1254	135
Average Education	11.5	0.99	6.66	13.0	135
Percentage Black	8.65	11.6	0	51.1	135
Percentage Farmer	10.5	9.2	0.3	57.1	135

Table B.2: MSA-Level Housing Regulations Regressions

Sample: Largest 135 MSAs			
	(1) WRLURI (2005)	(2) WRLURI (2005)	(3) WRLURI (2005)
Δ Fertility Rate (1940-1960)	0.755 (0.343)	0.490 (0.342)	0.434 (0.354)
Observations	135	135	135
Average Education		x	x
Average Income (1940)		x	x
Black Percentage			x
Farmer Percentage			x
Region FE	x	x	x

Table B.3: Relationship Between Wharton Residential Land Use Regulation Index (2005) and WWII Casualty Rates

Sample: Largest 135 MSAs

	(1) WRLURI (2005)	(2) WRLURI (2005)	(3) WRLURI (2005)
WWII Casualty Rate (%)	-4.822	-4.766	-4.684
	(2.271)	(2.185)	(2.430)
Observations	135	135	135
Average Education		x	x
Average Income (1940)		x	x
Black Percentage			x
Farmer Percentage			x
Region FE	x	x	x

Table B.4: 2SLS Estimates: Housing Regulations and Changes in Fertility (Instrumented with WWII Casualty Rates)

Sample: Largest 135 MSAs

	(1) WRLURI (2005)	(2) WRLURI (2005)	(3) WRLURI (2005)
Δ Fertility Rate (1940-1960)	1.659	1.509	1.531
	(0.709)	(0.704)	(0.804)
Observations	135	135	135
Average Education		x	x
Average Income (1940)		x	x
Black Percentage			x
Farmer Percentage			x
Region FE	x	x	x

APPENDIX C

Chapter 3 Appendix

C.1 One-Link Analysis: Rank-Rank Slope

Table C.1: Rank-Rank Slopes of Income and Sales Amount

Dependent Variable: Characteristics of Buyer

	(1) Percentile Household Income	(2) Percentile Household Income	(3) Percentile Median Tract Income	(4) Percentile Sale Amount	(5) Percentile Sale Amount
Percentile Median Tract Income (Seller)			0.388 (0.007)		
Percentile Household Income (Seller)	0.318 (0.004)	0.332 (0.007)			
Percentile Sale Amount (Seller)				0.519 (0.003)	0.583 (0.006)
Observations	69979	18891	18891	69979	18891
Sample	Full Sample	Linked HMDA	Linked HMDA	Full Sample	Linked HMDA

Standard errors in parentheses

C.2 Vacancy Chains: Transaction Dates of Match

Figure C.1: Link 1 - Date Gap Distribution: Sell Date Relative to House Purchase Date

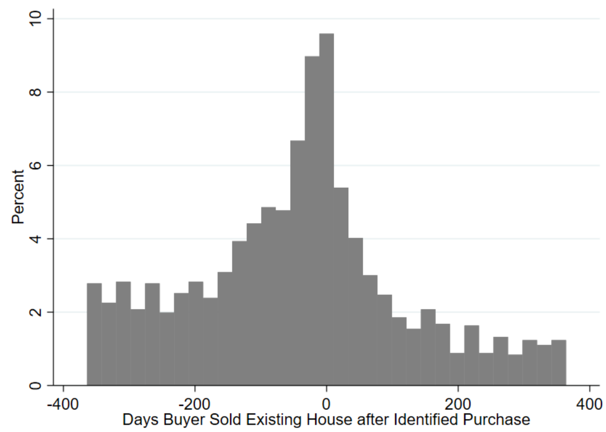

Figure C.2: Link 2 - Date Gap Distribution: Sell Date Relative to House Purchase Date

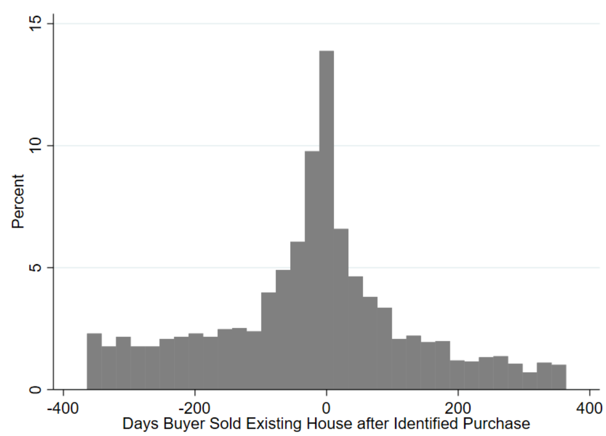

Figure C.3: Link 3 - Date Gap Distribution: Sell Date Relative to House Purchase Date

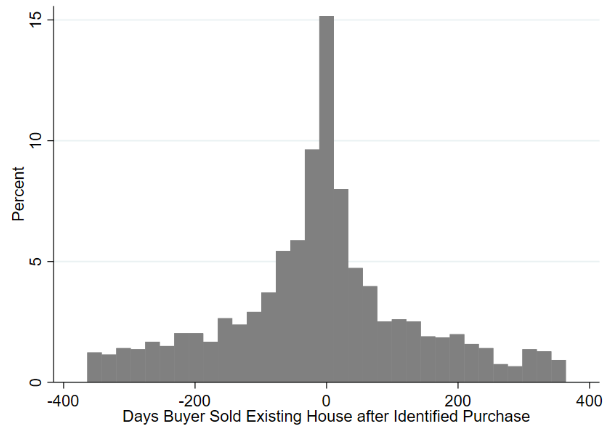

C.3 Vacancy Chain Analysis: Rank-Rank Slope

Table C.2: Rank-Rank Slope

Dependent Variable: Characteristics of Buyer

	(1) Percentile Median Tract Income	(2) Percentile Median Tract Income	(3) Percentile Household Income	(4) Percentile Sale Amount	(5) Percentile Sale Amount
Percentile Household Income (Seller)			0.299		
			(0.026)		
Percentile Median Tract Income (Seller)	0.259	0.263			
	(0.012)	(0.027)			
Percentile Sale Amount (Seller)				0.363	0.438
				(0.012)	(0.026)
Observations	6786	1162	1162	6786	1162
Sample	Full Chain	Linked HMDA	Linked HMDA	Full Chain	Linked HMDA

Standard errors in parentheses

Table C.3: Rank-Rank Slope of Income by Link

Dependent Variable: Tract, HH Income, or Sales Amount

	(1) Percentile Median Tract Income	(2) Percentile Median Tract Income	(3) Percentile Household Income
Percentile Median Tract Income (Seller)	0.236	0.280	
	(0.021)	(0.047)	
Percentile Median Tract Income (Seller) * Link=2	0.038	0.004	
	(0.029)	(0.066)	
Percentile Median Tract Income (Seller) * Link=3	0.030	-0.053	
	(0.029)	(0.067)	
Percentile Household Income (Seller)			0.281
			(0.048)
Percentile Household Income (Seller) * Link=2			0.020
			(0.067)
Percentile Household Income (Seller) * Link=3			0.032
			(0.065)
Observations	6786	1162	1162
Sample	Full Chain	Linked HMDA	Linked HMDA

Standard errors in parentheses

C.4 Alternative Visualization: Distribution of Characteristics by Link

For Median Tract Income and Sales Amount Below, the Link 0 distribution for each link is normalized to 100.

Figure C.4: Median Income (Tract) by Link

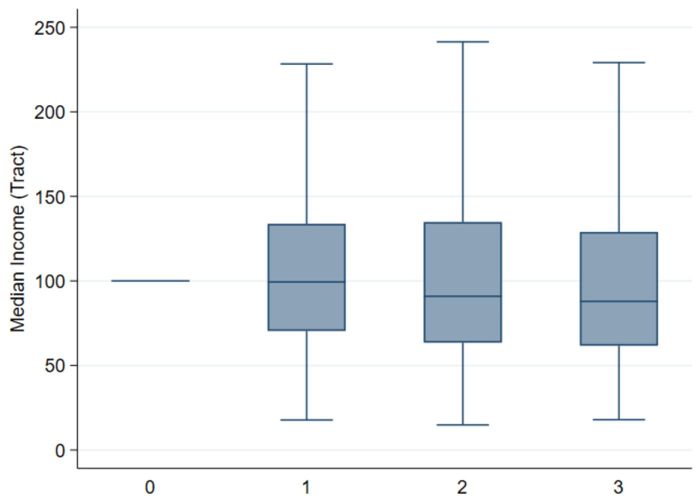

Figure C.5: Sales Amount by Link

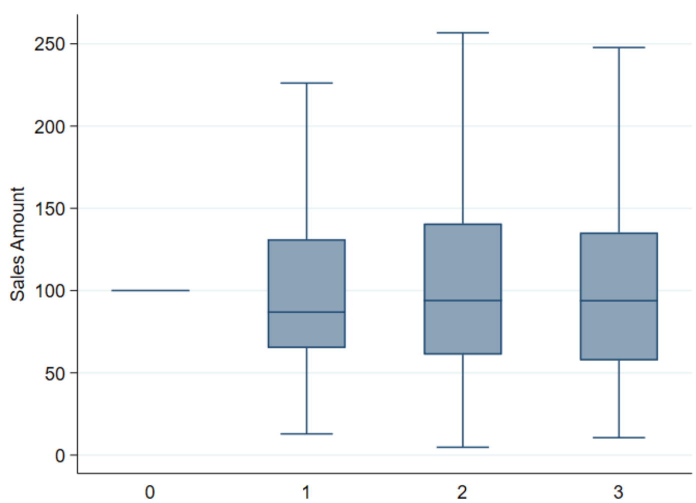

Figure C.6: Minority Population (Tract) by Link

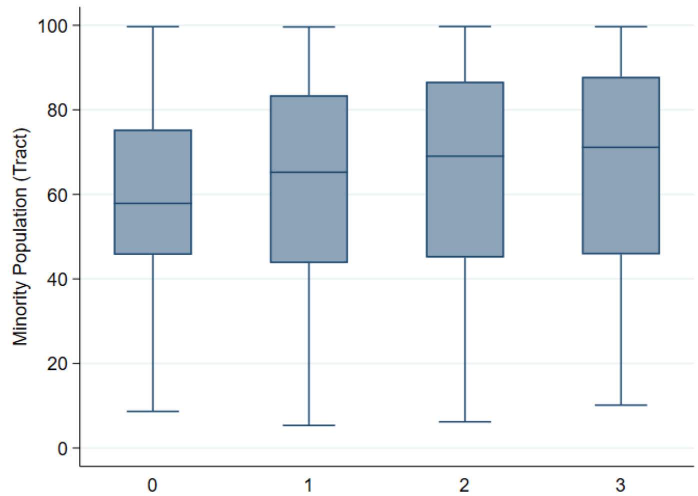

Figure C.7: Poverty Percentage (Tract) by Link

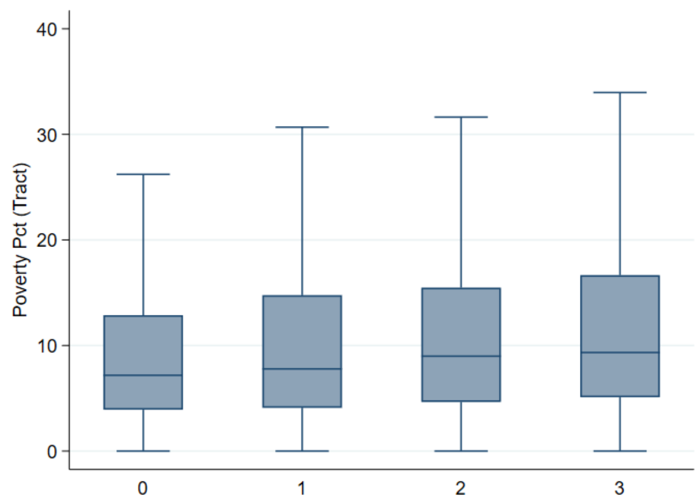

C.5 Vacancy Chain Analysis: Link Specific Transition Matrices

Table C.4: Vacancy Chain Transition Matrix (Link 1): Race

Buyer Race \ Seller Race	Non-Hispanic White	Hispanic White	Black	Other Minority
Non-Hispanic White	0.4789	0.2841	0.1667	0.2769
Hispanic White	0.2746	0.5227	0.5556	0.3231
Black	0.0211	0.0455	0.0555	0.0385
Other Minority	0.2584	0.2213	0.2435	0.3917

Table C.5: Vacancy Chain Transition Matrix (Link 2): Race

Buyer Race \ Seller Race	Non-Hispanic White	Hispanic White	Black	Other Minority
Non-Hispanic White	0.4709	0.2698	0.1667	0.3143
Hispanic White	0.2326	0.5001	0.5833	0.2000
Black	0.0058	0.0238	0.0000	0.0095
Other Minority	0.2907	0.2063	0.2500	0.4762

Table C.6: Vacancy Chain Transition Matrix (Link 3): Race

Buyer Race \ Seller Race	Non-Hispanic White	Hispanic White	Black	Other Minority
Non-Hispanic White	0.5772	0.2753	0.5714	0.2623
Hispanic White	0.1544	0.4928	0.0000	0.2049
Black	0.0268	0.0362	0.2857	0.0410
Other Minority	0.2416	0.1957	0.1429	0.4918

Table C.7: Vacancy Chain Transition Matrix (Link 1): Loan Type

$$\begin{array}{c c} & \text{Seller Loan Type} \\ & \begin{array}{cc} \text{Conventional} & \text{FHA/VA} \end{array} \\ \text{Buyer Loan Type} \quad \begin{array}{r} \text{Conventional} \\ \text{FHA/VA} \end{array} & \left[\begin{array}{cc} 0.7390 & 0.5566 \\ 0.2610 & 0.4434 \end{array}\right] \end{array}$$

Table C.8: Vacancy Chain Transition Matrix (Link 2): Loan Type

$$\begin{array}{c c} & \text{Seller Loan Type} \\ & \begin{array}{cc} \text{Conventional} & \text{FHA/VA} \end{array} \\ \text{Buyer Loan Type} \quad \begin{array}{r} \text{Conventional} \\ \text{FHA/VA} \end{array} & \left[\begin{array}{cc} 0.7629 & 0.5242 \\ 0.2371 & 0.4758 \end{array}\right] \end{array}$$

Table C.9: Vacancy Chain Transition Matrix (Link 3): Loan Type

$$\begin{array}{c c} & \text{Seller Loan Type} \\ & \begin{array}{cc} \text{Conventional} & \text{FHA/VA} \end{array} \\ \text{Buyer Loan Type} \quad \begin{array}{r} \text{Conventional} \\ \text{FHA/VA} \end{array} & \left[\begin{array}{cc} 0.6971 & 0.5070 \\ 0.3029 & 0.4930 \end{array}\right] \end{array}$$

BIBLIOGRAPHY

Abadie, Alberto (2021), "Using synthetic controls: Feasibility, data requirements, and methodological aspects." *Journal of Economic Literature*, 59, 391–425.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010), "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program." *Journal of the American Statistical Association*, 105, 493–505.

Acemoglu, Daron, David H. Autor, and David Lyle (2004), "Women, war, and wages: The effect of female labor supply on the wage structure at midcentury." *Journal of Political Economy*, 112, 497–551.

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst (2021), "Leisure luxuries and the labor supply of young men." *Journal of Political Economy*, 129, 337–382.

Albanesi, Stefania and Claudia Olivetti (2016), "Gender roles and medical progress." *Journal of Political Economy*, 124, 650–695.

Albouy, David and Gabriel Ehrlich (2018), "Housing productivity and the social cost of land-use restrictions." *Journal of Urban Economics*, 107, 101–120.

Albouy, David, Gabriel Ehrlich, and Yingyi Liu (2016), "Housing demand, cost-of-living inequality, and the affordability crisis." Technical report.

Anenberg, Elliot and Daniel Ringo (2022), "The propagation of demand shocks through housing markets." *American Economic Journal: Macroeconomics*, 14, 481–507.

Banks, James, Richard Blundell, Zoé Oldfield, and James Smith (2015), "House price volatility and the housing ladder." Technical report.

Becker, Gary Stanley and Kevin M. Murphy (2003), *Social Economics*. Belknap Press.

Bellet, Clement (2019), "The McMansion effect: Top size inequality, house satisfaction and home improvement in U.S. suburbs." *SSRN Electronic Journal*.

Bratu, Cristina, Oskari Harjunen, and Tuukka Saarimaa (2023), "JUE Insight: City-wide effects of new housing supply: Evidence from moving chains." *Journal of Urban Economics*, 103528.

Brodeur, Abel and Lamis Kattan (2022), "World War II, the baby boom, and employment: County-level evidence." *Journal of Labor Economics*, 40, 437–471.

Cavallo, Eduardo, Sebastian Galiani, Ilan Noy, and Juan Pantano (2013), "Catastrophic natural disasters and economic growth." *The Review of Economics and Statistics*, 95, 1549–1561.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez (2014), "Where is the land of opportunity? The geography of intergenerational mobility in the United States." *The Quarterly Journal of Economics*, 129, 1553–1623.

Deutscher, Nathan and Bhashkar Mazumder (2021), "Measuring intergenerational income mobility: A synthesis of approaches." Technical report.

Doepke, Matthias, Moshe Hazan, and Yishay D. Maoz (2015), "The baby boom and World War II: A macroeconomic analysis." *The Review of Economic Studies*, 82, 1031–1073.

Dubin, Frey A., D. Roderick Kiewiet, and Charles Noussair (1992), "Voting on growth control measures: Preferences and strategies." *Economics and Politics*, 4, 191–213.

Ferman, Bruno, Cristine Pinto, and Vitor Possebom (2020), "Cherry picking with synthetic controls." *Journal of Policy Analysis and Management*, 39, 510–532.

Ferrari, Ed (2011), "Conceptualising social housing within the wider housing market: A vacancy chain model." *Housing Studies*, 26, 95–116.

Fischel, William A. (2004), "An economic history of zoning and a cure for its exclusionary effects." *Urban Studies*, 41, 317–340.

Fischel, William A. (2005), *Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies*. Harvard University Press.

Ganong, Peter and Daniel Shoag (2017), "Why has regional income convergence in the U.S. declined?" *Journal of Urban Economics*, 102, 76–90.

Glaeser, Edward and Joseph Gyourko (2018), "The economic implications of housing supply." *Journal of Economic Perspectives*, 32, 3–30.

Goldin, Claudia and Claudia Olivetti (2013), "Shocking labor supply: A reassessment of the role of World War II on women's labor supply." *American Economic Review*, 103, 257–262.

Gray, M. Nolan and Adam A. Millsap (2020), "Subdividing the unzoned city: An analysis of the causes and effects of Houston's 1998 subdivision reform." *Journal of Planning Education and Research*.

Greenwood, Jeremy, Ananth Seshadri, and Guillaume Vandenbroucke (2005), "The baby boom and baby bust." *American Economic Review*, 95, 183–207.

Gyourko, Joseph and Raven Molloy (2015), "Regulation and housing supply." In *Handbook of Regional and Urban Economics*, 1289–1337, Elsevier.

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, 693–729.

Hamilton, Bruce W. (1975), “Zoning and property taxation in a system of local governments.” *Urban Studies*, 12, 205–211.

Hardman, Anna and Yannis M. Ioannides (2004), “Neighbors’ income distribution: economic segregation and mixing in US urban neighborhoods.” *Journal of Housing Economics*, 13, 368–382.

Hilber, Christian A.L. and Frédéric Robert-Nicoud (2013), “On the origins of land use regulations: Theory and evidence from US metro areas.” *Journal of Urban Economics*, 75, 29–43.

Hirt, Sonia A. (2015), *Zoned in the USA*. Cornell University Press.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider (2015), “The housing market (s) of San Diego.” *American Economic Review*, 105, 1371–1407.

Lansing, John B., Charles Wade Clifton, and James N. Morgan (1984), “New homes and poor people, 1966-1967.”

Levine, Jonathan (2005), *Zoned out*. Resources for the Future.

Lipsey, Robert E. and Doris Preston (1966), *Source Book of Statistics Relating to Construction*. National Bureau of Economic Research.

Mankiw, N.Gregory and David N. Weil (1989), “The baby boom, the baby bust, and the housing market.” *Regional Science and Urban Economics*, 19, 235–258.

Massey, D. S., G. A. Condran, and N. A. Denton (1987), “The effect of residential segregation on black social and economic well-being.” *Social Forces*, 66, 29–56.

Mast, Evan (2021), “JUE Insight: The effect of new market-rate housing construction on the low-income housing market.” *Journal of Urban Economics*, 103383.

McDonald, John F. (1995), “Houston remains unzoned.” *Land Economics*, 71, 137.

Mieszkowski, Peter and Edwin S Mills (1993), “The causes of metropolitan suburbanization.” *Journal of Economic Perspectives*, 7, 135–147.

Ortalo-Magné, François and Andrea Prat (2014), “On the political economy of urban growth: Homeownership versus affordability.” *American Economic Journal: Microeconomics*, 6, 154–181.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel (2020), “Segmented housing search.” *American Economic Review*, 110, 720–759.

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

Quigley, John M., Steven Raphael, and Larry A. Rosenthal (2008), “Measuring land-use regulations and their effects in the housing market.” *UC Berkeley: Berkeley Program on Housing and Urban Policy*.

Rosen, Richard (2011), "Competition in mortgage markets: The effect of lender type on loan characteristics." *Federal Reserve Bank of Chicago Economic Perspectives*, 35.

Saiz, Albert (2010), "The geographic determinants of housing supply." *Quarterly Journal of Economics*, 125, 1253–1296.

Schelling, Thomas C. (1971), "Dynamic models of segregation." *The Journal of Mathematical Sociology*, 1, 143–186.

Schuetz, Jenny (2009), "No renters in my suburban backyard: Land use regulation and rental housing." *Journal of Policy Analysis and Management*, 28, 296–320.

Shertzer, Allison, Tate Twinam, and Randall P. Walsh (2022), "Zoning and segregation in urban economic history." *Regional Science and Urban Economics*, 94, 103652.

Turner, Lena Magnusson (2008), "Who gets what and why? Vacancy chains in Stockholm's housing market." *European Journal of Housing Policy*, 8, 1–19.

White, Harrison C. (1971), "Multipliers, vacancy chains, and filtering in housing." *Journal of the American Institute of Planners*, 37, 88–94.

Whittemore, Andrew H. (2020), "Exclusionary zoning." *Journal of the American Planning Association*, 87, 167–180.

Winkler, William E. (1990), "String comparator metrics and enhanced decision rules in the Fellegi-Sunter model of record linkage." *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 354–359.

Zabel, Jeffrey and Maurice Dalton (2011), "The impact of minimum lot size regulations on house prices in eastern Massachusetts." *Regional Science and Urban Economics*, 41, 571–583.

Data Sources

- **Home Mortgage Application Act LAR Data**: Federal Financial Institutions Examination Council and the Consumer Finance Protection Bureau. 2007-2017.

  https://www.consumerfinance.gov/data-research/hmda/historic-data/

- **IPUMS Public Census Data**: Steven Ruggles, Sarah Flood, Sophia Foster, Ronald Goeken, Jose Pacas, Megan Schouweiler and Matthew Sobek. IPUMS USA: Version 11.0 [dataset]. Minneapolis, MN: IPUMS, 2021.

  https://doi.org/10.18128/D010.V11.0

- **IPUMS Public CPS Data**: Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren and Michael Westberry. Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [dataset]. Minneapolis, MN: IPUMS, 2021.

  https://doi.org/10.18128/D030.V9.0

- **Corelogic Tax and Deed Data**: Corelogic (R). These data files are licensed materials obtained with the assistance of the University of Michigan Libraries.

- **TIGER/Line Files** Census 2000 TIGER/Line Files [machine-readable data files]/prepared by the U.S. Census Bureau-Washington, DC; 2000.

- **Wharton Residential Land Use Regulation Index 2005:** Gyourko, J., Saiz A., and Summers A. (2008). A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index, Urban Studies, Vol. 45, Iss. 3, pp. 693-729

- **Wharton Residential Land Use Regulation Index 2018:** Gyourko, J., Hartley, J., & Krimmel, J. (2019). The Local Residential Land Use Regulatory Environment Across U.S. Housing Markets: Evidence from a New Wharton Index. (No. w26573). National Bureau of Economic Research.

- **World War II, the Baby Boom, and Employment: County-Level Evidence: Data File (Supplemental Materials: Data Files)** Abel Brodeur and Lamis Kattan (2022). Journal of Labor Economics, Volume 40, Number 2.

  https://doi.org/10.1086/715485