

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

[Link to Most Recent Version](#)

Mike Mei*

October 17, 2022

Abstract

What are the distributional effects of the minimum lot size (MLS) regulation on household welfare? An overlooked channel is how the MLS regulation increases physical house size. 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 an increase in the marginal cost of house size of around 14%. To quantify the distributional welfare effects stemming from these incentives, I build a quantitative model with housing and demographics and show that the effects of observed price changes induced by MLS regulations disproportionately hurt lower income and smaller households. Specifically, I find that the bottom decile of households (in terms of household size and income) are hurt about \$25,000 more (in 2010 dollars) than the top decile. 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.

JEL Codes: R21, R23, R28, R52

Keywords: Houston minimum lot size, house size demand, urban and real estate policy, land use policy, economic demography

*(Job Market Paper) University of Michigan, mikemei@umich.edu. This research was supported in part by an 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 is solely the responsibility of the author and does not necessarily represent the official views of the NICHD. I thank Hoyt Bleakley, Ana Reynoso, Gabriel Ehrlich, Jonathan Levine, John Bound, and John Leahy for advising on this paper. Also, this paper began in Martha Bailey's economic demography class. Thank you to seminar and conference participants at the University of Michigan, Yale University, and the 2022 Urban Economics Conference in DC for extremely helpful comments. Finally, thank you to graduate students at Michigan Economics for feedback and advice.

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 most direct effect is the bundling of additional land 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, which is the focus of this paper, is the increase in the optimal size of housing by making it cheaper on a marginal basis to build larger, since the land to build larger would have been already acquired under a large MLS regulation. Thus, perhaps counterintuitively, the marginal cost of an additional unit of size (say, square feet) decreases under a (sufficiently large) MLS regulation because the additional land or space needed to build housing is already required. Hence, an MLS regulation can increase demand for housing size by distorting the marginal cost of size to be lower than it would otherwise be, causing substitution effects into larger housing.


This additional overlooked channel of the MLS regulation is important primarily because it has 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 \$25000 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 its heterogeneity) due to these MLS regulations, as well as the channels in which the characteristics of the family can amplify these costs.

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.

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

Figure 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.

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)—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 housing regulations discourage living closer to central business districts, there are environmental and congestion externalities involved with transportation.

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 Constant Elasticity of Substitution Preferences) in order to capture many important features that are relevant for quantifying the heterogeneity in welfare costs: the curved 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 an 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 covenants 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 12.5 log points and increased the marginal cost of house size by about 14 log points.

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. The key parameters of the model are estimated from population Census data. The simulation inputs into the model are disciplined by an event study analysis of the Houston minimum lot size deregulation in 1998. I show the simulated welfare results on households and their distribution 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 Background and Literature

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) [13], Gyourko et al. (2008) [16], Albuoy and Ehrlich (2018) [5], and Ganong and Shoag (2017) [12] covering important measures and costs of housing regulations. Hirt (2015) [17] 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) [21] and Gray and Milsap (2020) [14]. However, the literature has not adequately covered the precise impact of housing regulations on house size demand nor the relevant marginal costs 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) [18] predicted a housing bust after the Baby Boom. Banks et al. (2015) [6] 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. 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) [9]], technological innovations in the household [Greenwood et al. (2005)] [15], and maternal health innovations [Albanesi and Olivetti (2016)] [4]. I do not take a particular stance on the underlying reason behind changes in family size; 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 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) [11]. 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 and matches the time series of increasing house size.

1.1.2 Housing Size

The history of increasingly larger homes in the United States is a rather complicated one. Hirt (2015) [17] write that part of the demand for larger homes come from deeply embedded preferences that are core to the notion of wild American frontier (which stand in contrast to European cities). This is in line with empirical research that support conspicuous consumption models of residential homes, as in Bellet (2019) [7].

However, much of the historical literature has discussed government policy as a large cause of larger homes in cities. This includes the multiple determinants of suburbanization, as discussed in Mieszkowski and Mills (1993) [19]. A more recent literature looks specifically at zoning policies, as in Schuetz (2009) [20], who finds that restrictive zoning policies likely decreased quantity of smaller rental housing built, but the overall effects on aggregate rent levels are unclear.

1.1.3 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).

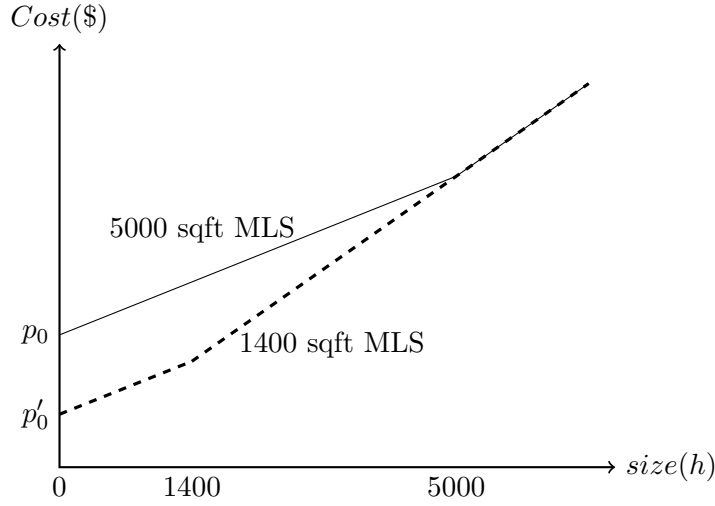
There is also substantial work in making sure the welfare calculations are realistic. With respect to heterogeneity in the minimum lot size's welfare costs over the income distribution, one of the most relevant factors is the shape of the Engel curve for house size. This is because housing is known to be a strong necessity good (which is a good where demand increases per unit of income, but at a diminishing rate). 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 is the housing setting 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 Section 4.1.1 about Model Fit.

2 Theoretical Framework

2.1 Basic Mechanism

Figure 2: Cost Structure Under Two Different Minimum Lot Sizes




The precise mechanism for how a minimum lot size regulation affects housing size and housing costs 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 price of an additional unit of house size is higher relative to a regime with a binding (say, 5000 square feet) minimum lot size, simply because more or less land is needed to build a certain size.

The following 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. Instead, the focus is on the main channels. 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. Hence, combined with the previous paragraph, the cost p to build a house of size h can be expressed in a reduced form way 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. The two scenarios, in which a 5000 square feet MLS is compared to a 1400 square feet MLS, are illustrated above.

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."

Figure 3: Change in MLS leads to Rotating Shift in Budget Constraint



Above, 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 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 cost 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).

2.1.1 Welfare Heterogeneity

In the simple model above illustrated in Figure 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 household 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.

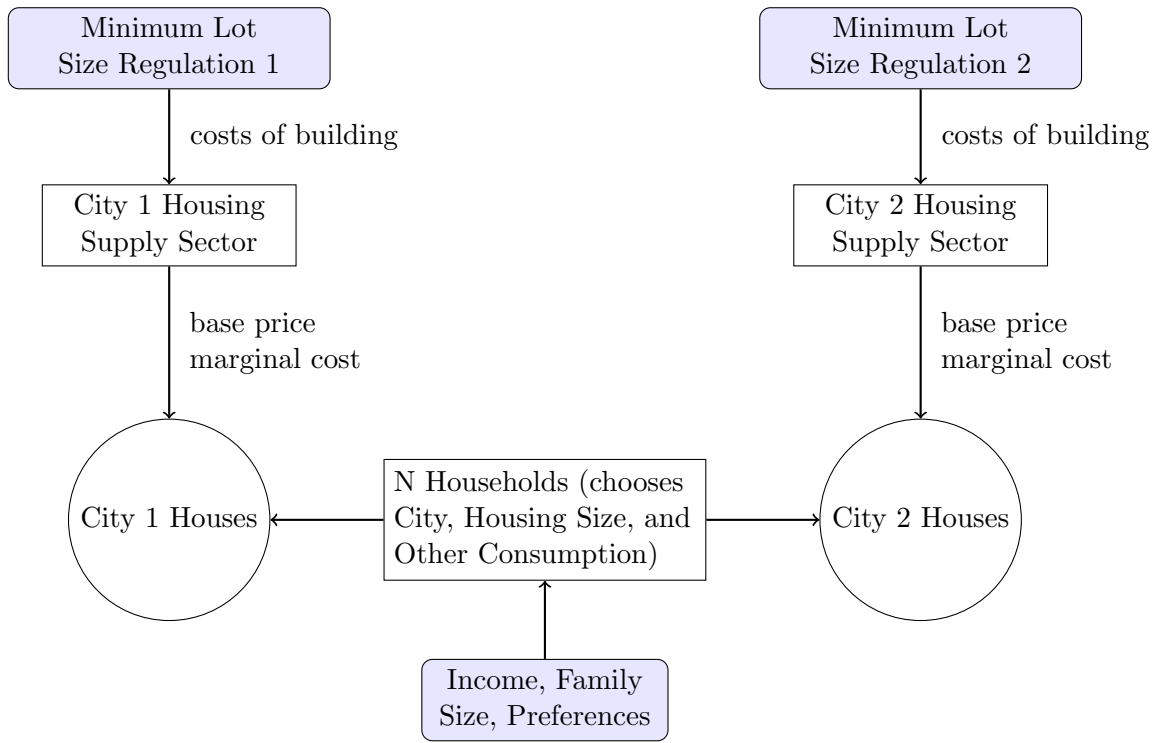
In essence, these two channels represent a decomposition of the total welfare gains into two parts. The first part is the mere difference between the the dotted line budget line and the solid budget line, representing the extra value of consumption if house size could not change; the nature of the MLS 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 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 positive correlated with income and household size.

2.2 Outline of Model Structure

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

Figure 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.

2.3 Housing Demand

Following the spirit of Aguiar et al. (2021) [3], the household decision 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. L represents location
2. θ_{it} is an age/family shifter of size demand
3. ξ_i, ζ_i^L are an idiosyncratic preference term for housing and location
4. η_c and η_h are elasticities

The budget constraint in each location L differs. Particularly:

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

where $p^L(h_t) = p_0^L + p_h^L h_t$ is the pricing function for housing in location L . Although it has a linear form, the estimation allows for nonlinear pricing in the form seen in Figure 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 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).

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

Second, as seen above, 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.² Specifically:

$$\theta_{AZ} = \alpha_0 + \alpha_1 A + \alpha_2 A^2 + \alpha_3 Z \quad (2)$$

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. (2017), 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

$$M_i = G[wY_i + (1 - w)\bar{Y}_g] \quad (3)$$

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

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

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

where

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

²Several 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.

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:

$$\underbrace{\log(\sigma\gamma) + (\gamma - 1) \log I_t}_{\text{common across types}} = \underbrace{\log PV_{qt} - \log \pi_{qt}}_{\text{type specific}} \quad (4)$$

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:

$$PV_q = \frac{1+r}{r+\delta} p_q \quad (5)$$

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.

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.

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 cost 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.

3 Empirical Analysis - Houston

To understand the effects of a minimum lot size regulation, I study a minimum lot size change in Houston that was enacted in 1998 and implemented in 1999. The seminal reference for this policy change is Gray and Milsap (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 vast majority of the population of Houston. For the rest of this paper, I define Houston to be only the areas within the I-610 loop, so any analysis of Houston excludes the parts of Houston outside of that loop.

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.³ Gray and Milsap (2020) has shown that this deregulation event spurred development of many smaller lots in middle income neighborhoods. Figure 5 shows a count of smaller lots (< 5000 square feet) developed. Because the Corelogic data has reliability problems before 1991, I augmented my graph with data from Gray and Milsap (2020):

Figure 5: Number of Smaller Lots (< 5000 SQFT) Developed in Houston



Note: Corelogic data counts augmented with Gray and Milsap (2020) series before 1991.

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 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

³Attempts to exclude the census tracts where there was the largest number of opt-outs to the policy did not change the results qualitatively.

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. In Figure 6, I show an example of a house that was built before the reform. Many single family houses then spanned a part of a large plot of land. A significant amount of the lot was used as backyard or landscaping.

In contrast, Figure 7 shows an example of post-reform housing. 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 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.

Figure 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 about 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 cost 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 study the effects on Houston.

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 1960 to 2010) and the American Community Survey (in years 2017) provided by IPUMS (see Data Sources). 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 dating back to 1960, 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.

3.2 Empirical Model: House Size

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 difference-in-difference framework, 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, Diamond, and Hainmueller (2010) [2] and subsequent work by Abadie (2021) [1]. 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 below.

Table 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.

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. In 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.

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 used follow the spirit of Abadie, Diamond, and Hainmueller (2010) in the use of a combination of evenly spaced out outcome variables and other predictor variables. This is a compromise between competing styles in the literature. For example, Ferman et al. (2020) [10] emphasize the importance of matching on a large number of pre-period treatment outcomes to avoid specification searching, while Cavallo et al. (2013) [8] practice limiting the matching to a few pre-period outcomes as a test of out-of-sample validity. Specifically, I use the odd-yearly 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 certain variables are dropped. In general, I show that the results are robust to choices of predictor variables.

Table 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 below. There are several 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 parts of the larger cities in Texas, which are natural and intuitive comparison cities to Houston.

3.3.2 House Size Results


In both figures below, 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

Table 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

its synthetic control. After 1999, however, the two series starts 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 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 9 plots the differences between Houston's trajectory and 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 statistic 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

Figure 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.

effects. Specifically, the test statistic for city i is:

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

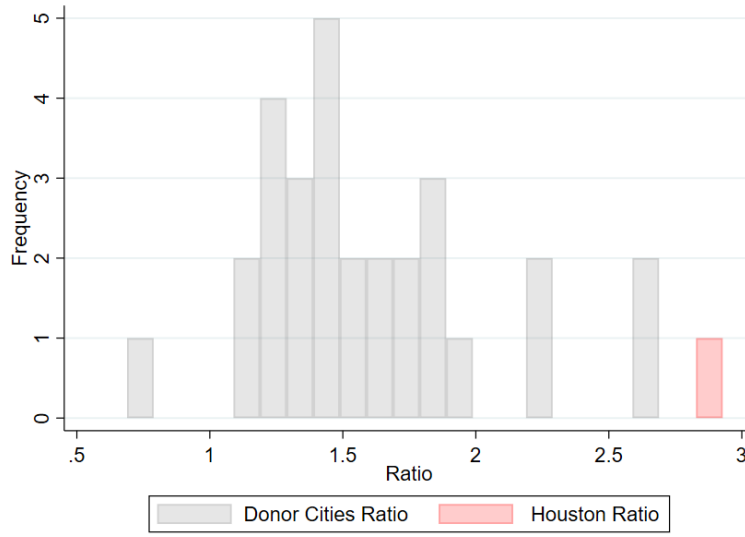
Intuitively, outcomes of cities that diverge significantly from its synthetic control, relative to that divergence before the treatment, show more statistically relevant results. The exact Fischer p-value is therefore the rank of this ratio (amongst all donor cities) as a fraction of the total number 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. Figure ?? shows the distribution of these ratios and visually shows that the Houston's test statistic ratio stands out.

For robustness checking 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).

3.3.3 Diff-in-diff Entire Distribution

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

Figure 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.

Figure 11: Synthetic Control Cities vs. Houston



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 marginal cost of house size.

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, 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 increased about 6 percentage. That presents about a 11 log point increase. The third regression

Table 4: Floor Area Ratios Increased in Houston

Dependent Variable: Different Measures of Floor Area Ratio			
	(1) FAR	(2) Log FAR	(3) $\mathbb{1}(FAR \geq 0.5)$
Houston*Post	0.060 (0.009)	0.116 (0.053)	0.048 (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			

specification suggests that a significant proportion of this rise was the shift to FARs higher than 0.5.

3.4 Empirical Model: Price of House Size

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 interacted 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_{icjt} 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. Our 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, I denote the pooled coefficients for the post period relative to the pre-period as β and α in the table below. I run specifications with different comparison groups (all donor cities vs. restricted to the synthetic control cities as identified in the previous section).

3.4.2 Price Results

Table 5: Price Regressions: Marginal Cost Increased in Houston

Dependent Variable: Sales Price (2010 dollars)			
	(1) Price	(2) Price	(3) Price
Houston*Post (α) (in thousands)	-27.721 (7.022)	-12.404 (3.866)	-11.924 (3.079)
Houston*Post*SQFT (β)	24.860 (6.195)	20.917 (3.827)	22.266 (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 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. In log terms relative to the marginal cost per square foot in the Houston pre-period, this is an **increase of about 14 log points** (i.e., $\Delta \log p \approx 0.14$). Moreover, the second specification with year effects (instead of pre/post effects) suggests that the overall base price of housing decreased about \$12000; these intercept estimates, however, are relatively noisy compared to the marginal cost (slope) estimates.

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.

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

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 are consistent with that. After the inclusion of lot size controls, the estimated change in the marginal cost 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.

There are several main takeaways from the empirical results from Houston. The first is that average house 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

Table 6: Controlling for Lot Size, Marginal Cost Increases Not As Significant

Dependent Variable: Sales Price (2010 dollars)			
	(1) Price	(2) Price	(3) Price
Houston*Post (α) (in thousands)	-28.122 (6.423)	-13.799 (5.068)	-9.691 (5.821)
Houston*Post*SQFT (β)	7.231 (2.237)	4.559 (1.789)	2.775 (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

cities). Finally, the ratio of identified house size and price effects imply a price elasticity of house size demand to be $\epsilon = \frac{\partial \log(h)}{\partial \log p} = \frac{12.5}{14} = 0.893$ which is used to verify the fit of the model in the next section.

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.

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 2000 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 as well as provide a graph of model fit below.

Below 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

Table 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

income converges to the group average over time.

The housing size parameter, demographic shifters, and variances are estimated from the cross-sectional variances, the details of which are given below. 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 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_n , 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 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 parameter α_3 being positive 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.

Figure 12: **Baseline Model:** Simulated Model Averages vs. Empirical Averages for Each Age-HHSize-Income Group



4.1.1 Model Fit

4.1.2 Targeted Model Fit

Next I discuss the quantitative fit. In Figure 12, I plot the model generated moments against the empirical moments estimated from the Census data, excluding the moment associated with variance.⁴ 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.


I further explain my contribution of using these preferences to capture the curved Engel curves of housing. By breaking up the model fit line into lower (below 20th percentile) vs medium (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 even a decrease in the degrees of freedom. What is noticeable in the fit graph above, however, is the systematic way in the model overestimates housing size demand for higher income groups and underestimates them for lower income groups. This lack of fit is not only systematic but it significantly changes the quality of the fit. 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.

4.1.3 Untargeted Model Fit

A stronger test of a model is whether it can explain patterns and features of reality that is not imposed upon 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

⁴The 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.

Figure 13: **Alternative Homothetic Model:** Simulated Model Averages vs. Empirical Averages for Each Age / HH Size / Income Group



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, only comes from estimates of η_c and η_h that are estimated by data that heuristically use information about 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 3.4.2. The latter estimated elasticity is estimated from an arguably 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.

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 above. 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 cost of size of building 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 foot and house size. This pivot point is 4.5 bedrooms, which covers almost all new housing that was built in the pre or post periods of the reform. 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 differing intensive margins of bedroom size available for choice for people in the data.

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? 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.

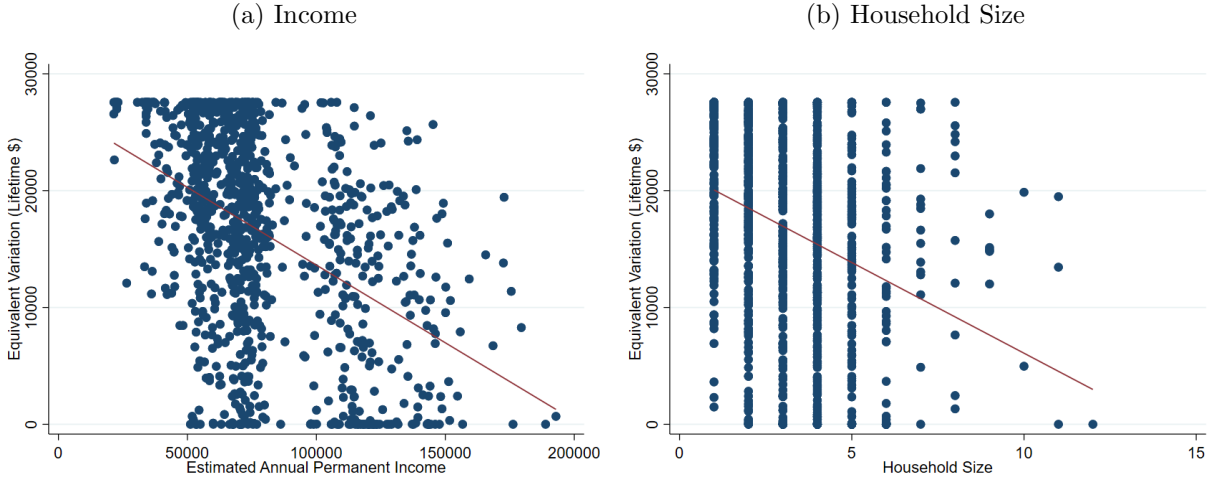
To induce the model to decrease housing size in the long run by 14 log points (baseline) or 3.8 log points (alternative), the marginal cost of housing size needs to rise by about 13 log points or 4.3 log points respectively. Since the estimated Houston drop in price was about 12.5 log points, this is 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.5, I show the results for the alternative specification, where the magnitudes of the effects are smaller, but the heterogeneity is still present.

4.3 Welfare Results

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 the demand for housing size as well as their differences in the willingness to re-optimize, which are all influenced 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 household 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 substantially

Figure 14: **Baseline Simulation:** Household Lifetime Gains Across Income and Household Size (2010 Dollars)



higher than the amount from the reduced form price regression. 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 \$25,000 between the households that have the most to gain (top decile in terms of household size and income) and those who have the least to gain (bottom decile).

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 many ways, this is also a calculation that takes into account 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 heterogeneity that is one magnitude larger. Similar to previous results, homeowners who originally had larger houses are hurt more by the policy change.

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

With the intuition derived in Section 2.1.1, the total welfare gains in Figure 14 can be decomposed into two parts: a part that comes from having to buy less lot size 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, but re-optimization is also economically important: (\$2,000 average gains). However, although re-optimization is increasing in income and household size, the heterogeneity is relatively small. Hence, the first order effects dominate and drive the direction of the heterogeneity in the baseline results.

Figure 15: **Baseline Simulation With Asset Price Effects:** Household Lifetime Gains Across Income and Household Size (2010 Dollars)

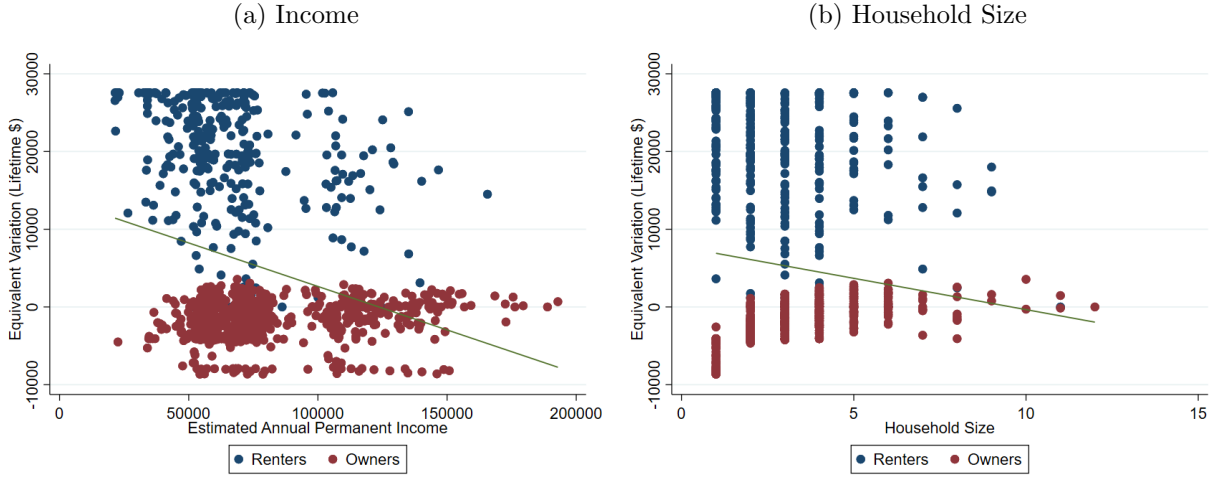
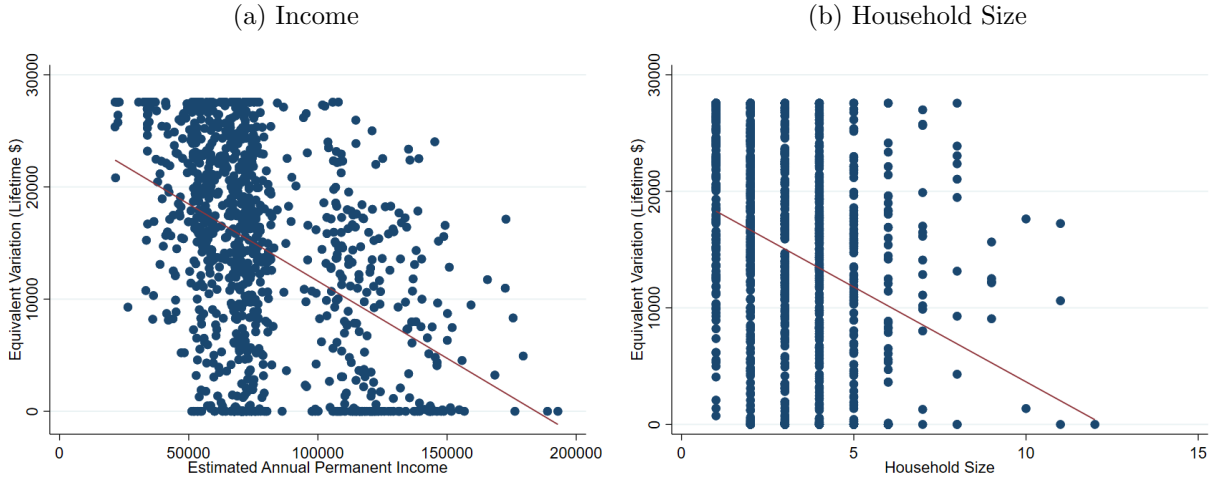


Figure 16: **Baseline Simulation:** Value of Lot Size Savings Across Income and Household Size (2010 Dollars)



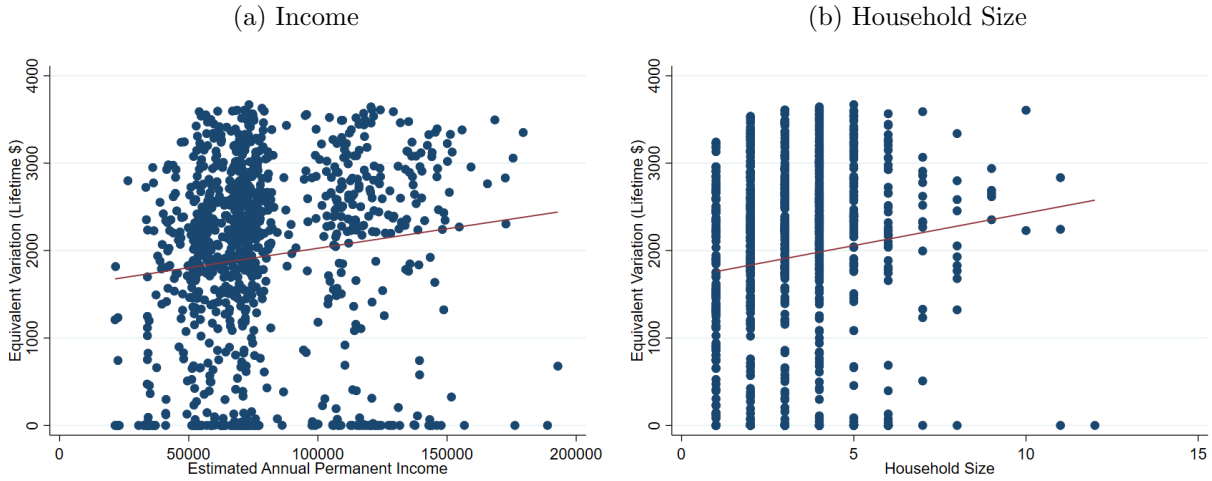
Notes: The value of lot size savings is simply the money saved given a household's initial house size demand, without the opportunity to re-optimize and choose a new house size.

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 cost. An additional important and natural prediction of this model is described in this section.

The model predicts 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 assumptions that influence the direction of any Roy Model selection include 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

Figure 17: **Baseline Simulation:** Re-optimization Gains Across Income and Household Size (2010 Dollars)



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.

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, rather than correlations in preferences across locations. For a further mathematical exposition of the Roy Model in this setting, see Appendix A.7 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 this theoretical result, I show the simulated result from the model. The table below shows the equivalent difference-in-difference estimator for selection on characteristics as predicted by the model with normal uncorrelated locational preference shocks.⁵ 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.

Table 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
$\bar{X}_{2,post}$	3.25	63353.4

⁵The 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.

I then turn to the data from the Current Population Survey. Although the geographies available is 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.

$$y_{ict} = \gamma_t + \lambda_c + \beta_{houston} * post + \epsilon_{ict} \quad (6)$$

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

Below are the results from the basic diff-in-diff 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 mere 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.

Table 10: Selection of People into Houston vs. Other Texas Cities

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

5 Conclusion

The general point of this paper is that there is 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 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 are 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.

References

- [1] Alberto Abadie. Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2):391–425, jun 2021.
- [2] Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American Statistical Association*, 105(490):493–505, jun 2010.
- [3] Mark Aguiar, Mark Bills, Kerwin Kofi Charles, and Erik Hurst. Leisure luxuries and the labor supply of young men. *Journal of Political Economy*, 129(2):337–382, feb 2021.
- [4] Stefania Albanesi and Claudia Olivetti. Gender roles and medical progress. *Journal of Political Economy*, 124(3):650–695, jun 2016.
- [5] David Albouy and Gabriel Ehrlich. Housing productivity and the social cost of land-use restrictions. *Journal of Urban Economics*, 107:101–120, sep 2018.
- [6] James Banks, Richard Blundell, Zoé Oldfield, and James Smith. House price volatility and the housing ladder. Technical report, jun 2015.
- [7] Clement Bellet. The McMansion effect: Top size inequality, house satisfaction and home improvement in u.s. suburbs. *SSRN Electronic Journal*, 2019.
- [8] Eduardo Cavallo, Sebastian Galiani, Ilan Noy, and Juan Pantano. Catastrophic natural disasters and economic growth. *The Review of Economics and Statistics*, 95(5):1549–1561, dec 2013.
- [9] Matthias Doepke, Moshe Hazan, and Yishay D. Maoz. The baby boom and world war II: A macroeconomic analysis. *The Review of Economic Studies*, 82(3):1031–1073, mar 2015.
- [10] Bruno Ferman, Cristine Pinto, and Vitor Possebom. Cherry picking with synthetic controls. *Journal of Policy Analysis and Management*, 39(2):510–532, mar 2020.
- [11] William A. Fischel. An economic history of zoning and a cure for its exclusionary effects. *Urban Studies*, 41(2):317–340, feb 2004.
- [12] Peter Ganong and Daniel Shoag. Why has regional income convergence in the u.s. declined? *Journal of Urban Economics*, 102:76–90, nov 2017.
- [13] Edward Glaeser and Joseph Gyourko. The economic implications of housing supply. *Journal of Economic Perspectives*, 32(1):3–30, feb 2018.
- [14] M. Nolan Gray and Adam A. Millsap. Subdividing the unzoned city: An analysis of the causes and effects of houston’s 1998 subdivision reform. *Journal of Planning Education and Research*, page 0739456X2093515, jul 2020.
- [15] Jeremy Greenwood, Ananth Seshadri, and Guillaume Vandenbroucke. The baby boom and baby bust. *American Economic Review*, 95(1):183–207, feb 2005.
- [16] Joseph Gyourko, Albert Saiz, and Anita Summers. A new measure of the local regulatory environment for housing markets: The wharton residential land use regulatory index. *Urban Studies*, 45(3):693–729, mar 2008.

- [17] Sonia A. Hirt. *Zoned in the USA*. Cornell University Press, aug 2015.
- [18] N.Gregory Mankiw and David N. Weil. The baby boom, the baby bust, and the housing market. *Regional Science and Urban Economics*, 19(2):235–258, may 1989.
- [19] Peter Mieszkowski and Edwin S Mills. The causes of metropolitan suburbanization. *Journal of Economic Perspectives*, 7(3):135–147, aug 1993.
- [20] Jenny Schuetz. No renters in my suburban backyard: Land use regulation and rental housing. *Journal of Policy Analysis and Management*, 28(2):296–320, dec 2009.
- [21] Jeffrey Zabel and Maurice Dalton. The impact of minimum lot size regulations on house prices in eastern massachusetts. *Regional Science and Urban Economics*, 41(6):571–583, nov 2011.

Data Sources

- **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.

A 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. This variable was chosen over the total rooms measure because bedrooms tends to have a higher correlation with the actual size of houses as measured by square feet in the Corelogic data. 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.

$$N_{ij} = \alpha_j + \beta HH_{size}_{ij} + \lambda X_{ij} + \epsilon_{ij}$$

where N is number of bedrooms, HH_{size} is household size, and X has 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 (0.010)	0.235 (0.010)	0.166 (0.008)	0.189 (0.010)	0.182 (0.010)	0.221 (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. (2017), 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.

Table A.2: Engel Elasticities

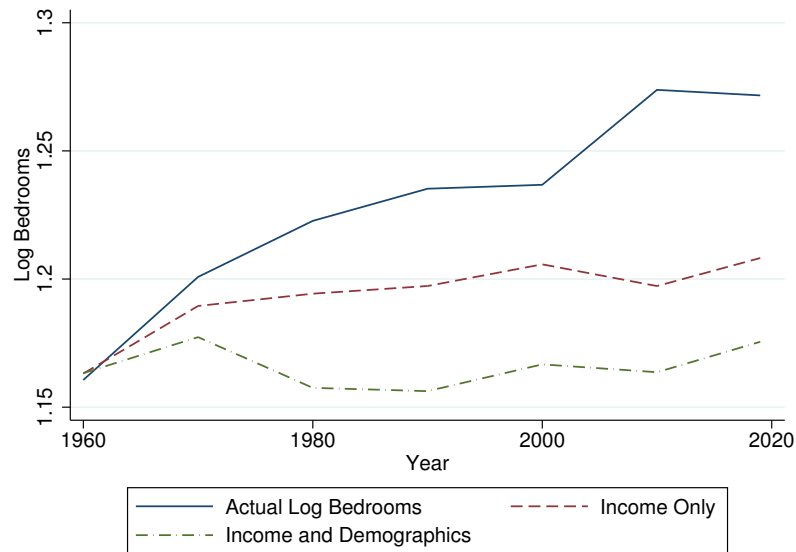
Dependent Variable: Log Number of Bedrooms				
	(1)	(2)	(3)	(4)
Log(Income)	0.102 (0.001)	0.110 (0.001)	0.072 (0.001)	0.077 (0.001)
Observations	482749	482749	480432	480432
Location FE		x		x
Demographic Controls			x	x
Standard errors in parentheses				

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).

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 would be similar.

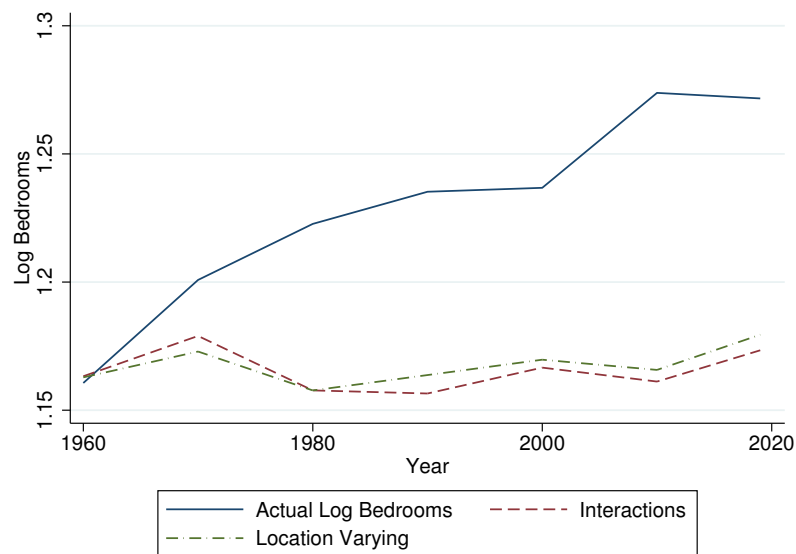
However, the following figures both show across a variety of specifications, the household side variables fall far short of explaining the historical trends. Even the specification with income only, which represents an unconditional Engel curve, shows that income by itself only explains about half of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

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>

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 the Appendix. 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. 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. In a subsequent model section, I will estimate prices from the Census data and disentangle those effects from shifts in technology/preferences over time. Finally, by looking at a hypothetical shock that affects only prices (through supply side effects) and not preferences, I can estimate the welfare changes and describe their heterogenous effects.

Table A.3: Engel Elasticities Alternative Estimates

Dependent Variable: Log Number of Bedrooms			
	(1) Baseline OLS	(2) Tobit	(3) IV
main			
Log(Income)	0.072 (0.001)	0.076 (0.001)	0.097 (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:

$$\begin{aligned} \max_{\{c_{it}\}, \{h_{it}\}} \quad & \sum_{t=0}^N \beta^t \left(\frac{c_{it}^{1-\frac{1}{\eta_c}}}{1-\frac{1}{\eta_c}} + \frac{(\theta_{it}\xi_i h_{it})^{1-\frac{1}{\eta_h}}}{1-\frac{1}{\eta_h}} \right) + \zeta_i^L \\ \text{s.t.} \quad & \sum_{t=0}^N \frac{c_{it} + p(h_{it})}{(1+r)^t} = M_i \end{aligned}$$

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:

$$(\beta(1+r))^t c_t^{\frac{-1}{\eta_c}} = \Omega \quad (7)$$

$$(\beta(1+r))^t (\theta_t \xi)^{1-\frac{1}{\eta_h}} h_t^{\frac{-1}{\eta_h}} = p_h \Omega \quad (8)$$

A.3.1 Relative Demand Equations and Numerical Solutions

Taking logs of the original first order condtions gives: Cancelling out the Ω terms gives the (Euler) intertemporal optimality conditions and the intratemporal optimality conditions below, respectively.

$$c_t^{\frac{1}{\eta_c}} = \beta(1+r) c_{t-1}^{\frac{1}{\eta_c}} \quad (9)$$

$$c_t^{\frac{1}{\eta_c}} = (\theta_t \xi)^{\frac{1}{\eta_h}-1} h_t^{\frac{1}{\eta_h}} p_h \quad (10)$$

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:

$$\frac{\log h_t}{\eta_h} - \frac{\log c_t}{\eta_c} = -\log p_h + \left(\frac{\eta_h - 1}{\eta_h} \right) \log \theta_t + \left(\frac{\eta_h - 1}{\eta_h} \right) \log \xi \quad (11)$$

A.3.2 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:

$$M_t = c_t + p_h h_t \quad (12)$$

$$= \Omega^{-\eta_c} (\beta(1+r))^{\eta_c t} + \Omega^{-\eta_h} (\theta_t \xi)^{\eta_h - 1} (\beta(1+r))^{\eta_h t} \quad (13)$$

Writing M_t and Ω in logs gives:

$$\log(M_t) = \log \left[e^{-\eta_c \log \Omega} (\beta(1+r))^{\eta_c t} + e^{-\eta_h \log \Omega} (\theta_t \xi)^{\eta_h - 1} (\beta(1+r))^{\eta_h t} \right] \quad (14)$$

Using the implicit function theorem on the equation above implies:

$$\frac{d \log \Omega}{d \log M_t} = \frac{-1}{-\left(\frac{-\eta_c \Omega^{-\eta_c} (\beta(1+r))^{\eta_c t} - \eta_h \Omega^{-\eta_h} (\theta_t \xi)^{\eta_h - 1} (\beta(1+r))^{\eta_h t}}{M_t} \right)} \quad (15)$$

$$= \frac{-1}{-\left(\frac{-\eta_c c_t - \eta_h p_h h_t}{M_t} \right)} \quad (16)$$

$$= \frac{-1}{\eta_c s_c + \eta_h s_h} \quad (17)$$

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:

$$\frac{d \log c_t}{d \log M_t} = \frac{d \log c_t}{d \log \Omega} \frac{d \log \Omega}{d \log M_t} \quad (18)$$

$$= \frac{\eta_c}{\eta_c s_c + \eta_h s_h} \quad (19)$$

$$= \frac{\eta_c}{\bar{\eta}} \quad (20)$$

An analogous derivation gives:

$$\frac{d \log h_t}{d \log M_t} = \frac{\eta_h}{\bar{\eta}} \quad (21)$$

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 magnitude 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.

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.

Lastly, the calculated exact p-value does not change because Houston's ratio of its post-period deviation from 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 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.

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

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.5 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)

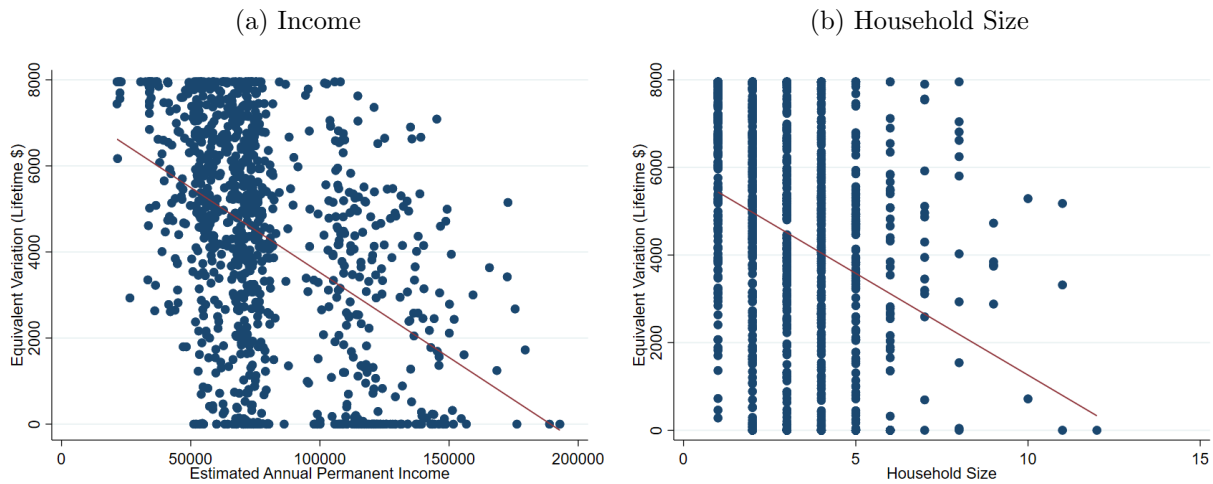
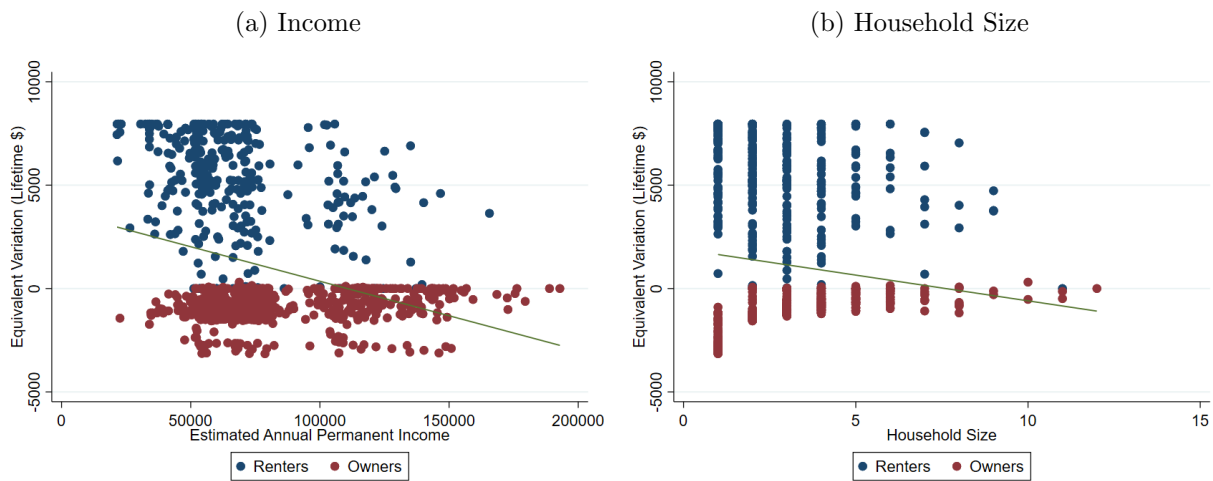


Figure A.7: **Alternative Simulation With Asset Price Effects:** Household Lifetime Gains Across Income and Household Size (2010 Dollars)



A.6 Robustness of Model to Choice of Calibrated Parameters


This section analyzes how the fundamental heterogeneity results (as well as overall model fit) compares 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.⁶ 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 \$4000 to \$7000. The largest discrepancy for the heterogeneity results comes from perturbing the discount factor term; this means that heterogeneity in welfare gains could be somewhat affected by how impatient households are. In the extreme limit, when households are so impatient that they only care about the present, the heterogeneity may be very small because only differences between households of the first periods' gains (i.e, when they are young) are relevant. Finally, the results do not seem to be qualitatively affected by large changes in the way that permanent income is weighted (between current income and cohort income), suggesting that measurement error on this margin may not affect the baseline results in first order terms.


⁶The idiosyncratic shocks are saved (or seeded, in computer science terms) and 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.

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) $w = 0.25$


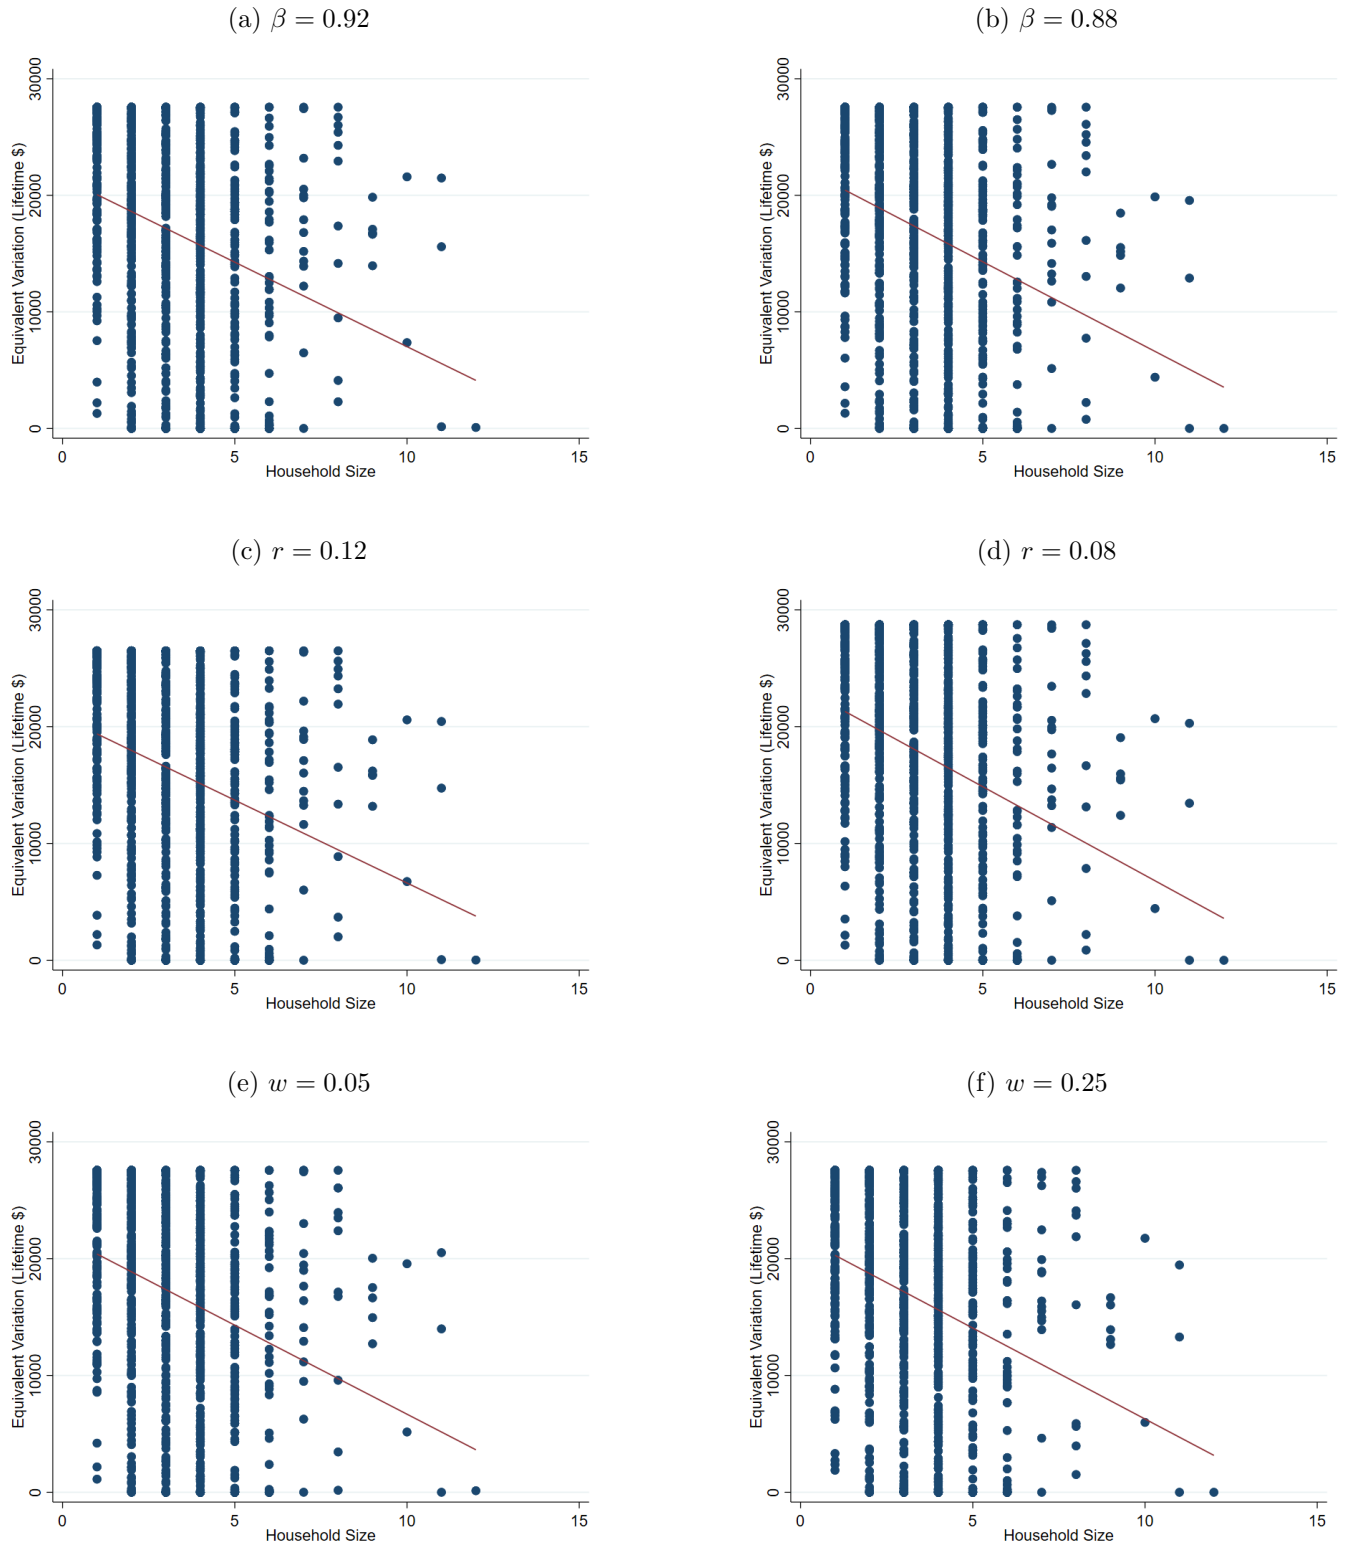


Figure A.9: Robustness: Welfare Heterogeneity Results Across Family Size, For Different Calibrated Parameters



A.7 Roy Model Sorting Mechanisms

A.7.1 The Canonical Roy Model

The Canonical Roy Model ⁷ 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 $\rho \gg 0$, 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.7.2 Housing Size Roy Model

In terms of the model in this paper, the destination location 1 is the city which reduces its minimum lot size. Source location is represented by location 0 which is not the city with the policy change.

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

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

where p^L are the vector of prices in each location, θ_i are a vector 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:

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

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

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

Taking the difference of the two equations:

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

A fundamental assumption consistent with this paper is that decreases in the price of size decreases the returns to additional income or household size. This is the underlying heterogeneity channel that affects

⁷“Lecture Note: Self Selection – The Roy Model.” David Autor. MIT. Accessed January 18, 2022. <https://economics.mit.edu/files/551>

the direction of selection. That is, assume:

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

for any variable F in ω .

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

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.