Estimating the Economic Value of Zoning Reform*

Santosh Anagol   Fernando Ferreira   Jonah Rexer

June 14, 2024

Abstract

We estimate the economic value of zoning reform in São Paulo, which altered maximum permitted construction along transportation corridors. Developers increased filings for multi-family construction in blocks affected by the reform, leading to more housing supply and lower housing prices in neighborhoods that allowed more densification. Our equilibrium model of housing markets estimates an aggregate 1.9% increase in housing stock and a 0.5% reduction in prices, resulting in large housing wealth transfers from current to future homeowners. The reform produced welfare gains of 0.76% of city GDP, mostly due to developer profits and consumer gains from the newly built environment.

Keywords: Housing supply, building restrictions, zoning reform, welfare gains, housing wealth.

JEL Codes: R0, H7, L8, K2.

*Santosh Anagol: anagol@wharton.upenn.edu, Wharton School, University of Pennsylvania. Fernando Ferreira: fferreir@wharton.upenn.edu, Wharton School, University of Pennsylvania and NBER. Jonah Rexer: jrexer@princeton.edu, School of Public and International Affairs, Princeton University. We are grateful for support from the Wharton Dean's Research Fund, and the Research Sponsors Program of the Zell/Lurie Real Estate Center. We thank Sherrie Cheng, Tom Cui, Anna Gao, Lucas Fernandes, Gi Kim, Renan Muta, Sophia Winston, Alexandru Zanca, and Holly Zhang for excellent research assistance. We also thank Rohan Ganduri, Edward Glaeser, Joseph Gyourko, Raven Molloy, and seminar participants at Imperial College Business School, Wharton Urban Lunch, European Urban Economics Association meeting, American Urban Economics Association meeting, NBER Summer Institute Real Estate, FGV-SP, Brazilian Society of Econometrics, LACEA-LAMES, NBER Public Economics, NTA meeting, Biennial Atlanta-Dallas Fed Conference, and the SITE conference at Stanford for helpful comments and suggestions. Finally, we thank Danilo Igliori and the DataZap team for providing the real estate listing data used in this research.

1 Introduction

The global urban population has been steadily increasing, from 33% in 1960 to 56% in 2019, with a predicted increase to 68% by 2050.[1] However, there is considerable uncertainty and debate regarding the ability of cities to provide adequate housing for their residents in the future. Satisfying the demand to live and work in cities will only be possible with a strong supply response from developers who build residential structures. In most cities, however, the supply of buildings is highly regulated by local authorities - the so-called "not in my back yard" NIMBYism - who can legally determine if and how developers can build (Glaeser, Gyourko and Saks, 2005; Gyourko and Molloy, 2015; Gyourko, Hartley and Krimmel, 2021).

Researchers attempting to assess the impact of zoning policies face four major issues. First, zoning policies are often endogenous, determined by price levels (e.g. high price areas choosing to restrict building to maintain high prices) – so comparisons across different jurisdictions can plausibly conflate price effects with zoning reform effects. Second, zoning policies are multi-dimensional and hard to measure and compare across time and locations; for example cities can regulate land use, building density, lot size, height, footprint, and setbacks, among other characteristics.[2] Third, there is a dearth of micro data on the behavior of developers, which is critical to estimating supply responses to local land use regulations. Fourth, even when these three issues are solved, it is quite challenging to trace out the impacts of zoning on built environments, price affordability, and welfare given the many aspects of supply and demand for housing impacted by local regulation. For example, new housing construction may improve the quality of the housing stock relative to the old and depreciated stock of homes, reduce housing prices, and allow households to move from the suburbs to the city and closer to their workplaces. But densification may also lead to increases in congestion, construction of less desirable apartment units (as opposed to single family housing), and loss of housing wealth by current homeowners and landlords, which is sometimes feared by proponents of NIMBYism.

Our paper shows how zoning reform and detailed data can be leveraged to address these four issues. We study a major reform in São Paulo, Brazil that generated granular, block-by-block changes in zoning parameters. The setting allows for a within-city approach that exploits both variation over time in zoning policy changes and block-level discontinuities in reform parame-

[1] United Nations Population Division. World Urbanization Prospects: 2018 Revision.

[2] The highly used Wharton Land Use Regulation Index (WLURI) is a summary measure of many land use restrictions collected via survey data. While WLURI is a useful proxy for overall land use regulation, it is difficult to translate causal effects of this index in to specific administrative land use parameters that might actually be reformed.

ters. We collect detailed administrative data to observe exactly which zoning parameters change and how these changes impact housing supply across the city at fine spatial resolution. By combining a reform with data on permitting activity and housing listings, we can credibly estimate the short-run responses of developers when land-use restrictions change, and the medium-run changes in housing supply and listing prices. Finally, we estimate an equilibrium model of housing supply and neighborhood demand to separately identify how changes in prices, quantities, and neighborhood amenities affect household welfare.

We examine the impact of a 2016 zoning reform in São Paulo, Brazil - the world's 4th largest metropolitan area with 21 million residents. The block-level reform centralized the ability to set density parameters previously under control of neighborhoods, and had the general goal of providing more dignified housing for its residents and allowing more densification along transportation corridors.[3] Each city block was assigned a maximum built-area-ratio (BAR) - the ratio of constructed square meters per square meter of land area - which defines the density of units that developers could develop on a given land parcel.[4] On average, the max BAR in the city's approximately 45,000 blocks increased from 1.54 to 2.09, allowing 36% more construction for a given lot size, and 45% of the city blocks had a max BAR increase of 1 or more. This unique variation in zoning restrictions over space and time allows us to control more precisely for the endogenous determinants of zoning policies.

We use a boundary discontinuity design to estimate developers' supply response to block-level changes in BAR. Blocks are categorized as treatment (where max BAR increased) or control (where max BAR stayed constant or decreased), and our running variable is the distance to the nearest block of the opposite category. Our identification exploits both cross-sectional and temporal changes in zoning regulation at the boundary discontinuity (in the spirit of a "difference-in-discontinuities" design).

We first apply this research design to rich administrative microdata on building permits. We find that the 1.4-point max BAR increase at the cut-off causes an increase of .003 multi-family

[3] The São Paulo reform's focus on increasing housing density near transportation corridors is similar to a recently failed reform aiming to increase housing density along California's public transit system. Mumbai's 2034 development plan also originally included, but then removed, plans for allowing high density construction near metro and commuter rail stations.

[4] We use the term BAR as it corresponds directly with the term used in Sao Paulo for this concept. It is closely related to the "FAR" ("floor-to-area"), "FSI" ("floor-space index") and "FSR" ("floor-space-ratio") metrics used in other parts of the world. See Brueckner and Singh (2020) for recent measurements for "FAR" stringency in five major U.S. cities. Sao Paulo has a unique set of rules defining what counts as constructed area; it includes livable space but excludes hallways, elevator shafts and in some cases mezzanine lobbies and swimming pools. Given these idiosyncracies we prefer to use the "BAR" terminology.

permits per block per quarter, corresponding to a sixty six percent increase in permits filed by developers. Differences in permitting activity between treatment and control blocks are absent prior to the zoning reform, and only start to emerge one year after the zoning reform was passed and increase thereafter, which is plausible given the time it likely takes for developers to create project plans, acquire land, and so on. We find no effect of a max BAR increase on single-family home permits filed or approved, which is consistent with the zoning parameter change easing constraints primarily for larger structures. We also estimate a spatial spillover model (following a strategy similar to Diamond and McQuade (2019)), finding no evidence of treatment effects on control blocks near the boundary discontinuity, suggesting minimal substitution of projects from control to treatment blocks. We find similar sized effects for treatment blocks near and far from the boundary, suggesting muted agglomeration effects so far.

We also estimate medium-run reform effects on housing market outcomes six years after the reform was enacted. For this exercise we use data from DataZap, a private Brazilian website similar to Zillow, that provides complete listings of apartments and houses for sale and their prices. We find that total housing units for sale increase by 10 percent in blocks that allow for more densification. The border discontinuity design, however, cannot assess price effects of zoning reforms since neighborhood amenities are generally not segmented by block.[5] We instead study neighborhood-level effects by aggregating availability of houses for sale and listing prices at the commuting zone level, which are cohesive neighborhoods with on average roughly 10,000 households. We find that neighborhoods with larger positive changes in max BAR have more growth in housing supply, and the effect is quantitatively similar to the block level boundary estimates. That same increase in max BAR results in listing price *declines*, indicating that the supply impact on prices dominates any price gains from improved amenities due to densification.

We then estimate an equilibrium model of supply and demand to predict the welfare consequences of the Sao Paulo 2016 reform. Our supply model extends the reduced form approach by adding the impact of prices, neighborhood features, and other regulatory zoning constraints. We estimate this model at the neighborhood level using Poisson regression to avoid the problem introduced by log-linear models when many locations have zero permits. Prices and zoning regulation in the form of max BAR are both endogenous. We instrument for BAR constraints using the boundary discontinuity-based variation, while at the same time controlling for other regu-

[5] In particular, Turner, Haughwout and Van Der Klaauw (2014) show that boundary discontinuity designs alone (including ours), cannot identify price, congestion and broader welfare effects which we would expect to affect both sides around a border experiencing a zoning change.

lations, such as rules limiting the size of a building's footprint (the "shadow ratio"), minimum setbacks, and height limits. We instrument for prices using localized demand-side Bartik shocks as in Baum-Snow and Han (2024). We find that a max BAR increase of 1 roughly doubles the number of permits, similar to our block-level boundary discontinuity model. A price increase of 20% generates a 50% increase in permits.

On the demand side, individuals maximize utility when choosing neighborhoods as a function of prices, location features, access to employment opportunities, and commuting time, allowing for residents to value both the positive and negative aspects of more densification. To estimate the model, we use survey data on 24,800 São Paulo metropolitan statistical area households which includes information on demographics, place of residence, and place of work. We instrument endogenous listing prices with geographic features in a ring outside the commuting zone; this strategy follows Berry, Levinsohn and Pakes (1995) and rests on the assumption that these outside features impact local prices through competition but do not directly impact the utility of individuals living in a given zone. To estimate consumers' preference for access to high paying jobs, we follow Tsivanidis (2022) and define an index of residential commuter access (RCMA) which measures a location's average travel time to zones with high paying jobs. Travel time varies according to neighborhood congestion. Finally, utility is driven by other important neighborhood features, such as age of buildings, number of units per building, neighborhood density, access to paved roads, average income, and the share of adults with a college degree.

Preferences for all of these neighborhood features are allowed to vary according to an individual's demographic characteristics. We model preferences using a multinomial logit framework and estimate the parameters with a two-stage maximum likelihood procedure, as in Bayer, Ferreira and McMillan (2007). We find negative elasticities of demand with respect to both prices and travel times, and a positive elasticity for RCMA; conditional on price and travel times, households prefer to live near areas with more high paying jobs. Consumers dislike old housing stock and greater density within a building, while valuing density in the broader neighborhood.

We then simulate welfare and distributional consequences for 2026 (ten years after the reform) under alternative zoning scenarios. We focus on 10 year forecasts as this gives sufficient time for permits to be converted to actual housing units, and for prices to adjust. The baseline scenario takes zoning parameters from the pre-2016 period, while in the second scenario we impose the new 2016 zoning map on our supply equation which leads to a supply shock generating new housing supply in different areas throughout the city. We then apply a market clearing condition that households' aggregate choice probabilities must equal supply shares for each neighborhood, which delivers an equilibrium price vector. To estimate welfare effects, we calculate neighborhood characteristics for each equilibrium. First, we use equilibrium changes in population to update average commuting times, which may rise after upzoning due to the congestion costs. Next, we use equilibrium quantities to update built environment characteristics, including building age and density. Finally, we evaluate households' utility functions at the new equilibrium to estimate the 2026 change in consumer surplus from the expected utility of access to neighborhoods with lower prices and new built environment characteristics.[6]

Counterfactual simulations show that the 2016 zoning reform leads to a 1.9% net increase in housing stock by 2026, resulting in a 0.5% average price reduction and modest welfare gains. However, there is variation across neighborhoods, with areas experiencing larger BAR shocks seeing more construction and greater price decreases, up to 17.1% and 4.6%, respectively. Such results are qualitatively similar to the reduced form estimates based on the medium-term listing data.

We then translate those equilibrium estimates into welfare gains. Average household consumer surplus increases by R$107 (approximately US$21), with college-educated and higher-income households benefitting the most from the reform, in part because they are the households most likely to take advantage of new housing. The average gains are mostly due to changes in the built environment, and the economic value of zoning reform is strongly linked to newer housing and preferences for densification. Congestion costs only slightly reduce welfare gains via increased commuting times.

Aggregating those estimates, we find that all consumer gains correspond to 0.08% of the city GDP. But the reform also impacted producers given the increased ability for developers to build new housing. Those extra developer profits amount to 0.57% of GDP. We also add a back-of-the envelope calculation of the productivity increases due to more agglomeration based on Glaeser and Gyourko (2018). This extra gain is 0.11% of GDP. Overall, total welfare gain from the 2016 zoning reform is 0.76% of the city GDP.

Finally, we also find that the housing price reductions impacted the nominal wealth of current homeowners and landlords. These nominal price changes are transfers from current to future real estate owners, and correspond to 1.88% of the city GDP. Such large transfers do not directly impact

[6] Our preferred model allows for the city population to grow as more housing is built, but we also test a model that assumes a closed economy. Details are shown in Section V.

overall city welfare. The main intuition for this result is that, if the price of housing declines due to a zoning reform, households who are selling their homes are made worse off but other households who are buying those same less expensive houses are made better off, offsetting welfare effects (Bajari, Benkard and Krainer, 2005). However, these wealth transfers may explain the lack of support for more dense construction on the part of current homeowners.

Related literature.: Zoning restrictions, such as limits on the density of buildings, are generally associated with increases in the cost of living (Glaeser and Gyourko, 2018; Brueckner and Sridhar, 2012; Ding, 2013), greater segregation and reduction in economic convergence (Trounstine, 2018; Ganong and Shoag, 2017), and result in the loss of economic output (Hsieh and Moretti, 2019). The literature studying the causal effects of zoning has primarily focused on cross-metropolitan area studies because of the variation in zoning rules across those geographies, with recent research focusing on block level differences around municipal boundaries (Turner, Haughwout and Van Der Klaauw (2014); Song (2021); Kulka, Sood and Chiumenti (2022);). Our main contribution here is to analyze how a zoning reform can be combined with micro-data to understand how zoning affects housing supply and long-run welfare.

Our paper also contributes to the literature on housing supply. One strand of this literature studies geographic constraints on housing supply (Saiz, 2010; Baum-Snow and Han, 2024). There is also a growing literature on structural models of housing supply. Murphy (2018) estimates a model of housing supply with a focus on the role of construction costs, while Paciorek (2013) studies the relationship between supply constraints and price volatility. Calder-Wang (2022) estimates the welfare of New York City residents given changes in availability of rental units due to the expansion of Airbnb.

Another literature focuses on understanding the internal structure of cities, with a focus on commuting and transit accessibility. Harari (2020) investigates how city shapes in India affect transit accessibility, land use regulations, and city growth. Ahlfeldt et al. (2015) estimate a model of internal city structure to quantify the effect of densification after the Berlin division and reunification. Tsivanidis (2022) estimates how new transit lines impact worker sorting and welfare in Bogota, and Balboni et al. (2020) estimates the impact of a new bus rapid transit system in Dar Es Salaam. We contribute to this literature by endogenizing housing regulations, and estimating how a zoning reform can affect availability of house units, affordability, and impact welfare.

This paper also fits in to three other broad topics. First, there is a large literature on the economics of urban density and agglomeration effects recently summarized by Duranton and Puga

(2020). Those authors note that future progress in this literature could encompass the dynamics of building construction and raise empirical standards in the identification of causal effects, both accomplished in our work. Second, there is a growing literature studying how housing supply will respond in the face of larger demand for cities, recently summarized in Brueckner and Lall (2015); our paper provides a first complete evaluation of a zoning reform in a developing country city context. Finally, Epple, Gordon and Sieg (2010) and Combes, Duranton and Gobillon (2021) directly estimate production functions for housing. Our paper contributes to that body of work by estimating in detail how developers respond to zoning reforms at very granular geographies, credibly dealing with the endogeneity of zoning parameters.

2 History of Zoning in São Paulo

São Paulo has had three major zoning reforms in the past fifty years: 1972, 2004 and 2016. All of these reforms created "zone types" such as "Mixed Residential Use," with each zone type assigned a set of building parameters. Each of the city's blocks are assigned to a zone type, and the zone type designation determines the block's building parameters. The primary building parameter set in each reform is the built-to-area ratio (BAR), which is the ratio between the computable area of the building and the lot size, and fundamentally determines density of units developers can build given a land parcel. Each block is also assigned to a building usage, such as residential, commercial or industrial purposes. Most zone types allow multiple building usages, but there are some that require only certain building types.

São Paulo's first city wide zoning regime was established in 1972 in response to rapid, haphazard, urban growth. Poor areas lacked enforcement of building rules allowing developers to build at their own will, but wealthy neighborhoods had rigorous regulation implemented by private developers. The 1972 zoning law primarily aimed to preserve the architecture of richer neighborhoods, while guiding the city's growth towards the periphery; this law also established that no new building could have a BAR above 4. However, the majority of urban land resided in zones with a max BAR of one.

São Paulo enacted a new Zoning Law in 2004, primarily in response to new federal laws mandating cities to have urban master plans. The 2004 Zoning Law had two features relevant for our analysis. First, the determination of specific building parameters, such as BAR, was decentralized to the "subprefeitura" or neighborhood level (São Paulo has 32 subprefeituras). In particular, the city would determine the zone type of each block, but the same zone type could have different building parameters based on the block's subprefeitura. Second, the city expanded the BAR parameter scheme to include minimum BAR, basic BAR and max BAR levels, each of which were chosen by the subprefeitura government for each zone-type within the subprefeitura. [7] The 2004 zoning reform also legislated, at the city level, that each district (a smaller unit than a subprefeitura with 96 total in the Sao Paulo municipality) - would have a limited stock of square meters above the basic BAR that could be constructed. Once a given district exhausted its available capacity above the basic BAR, no further construction could occur above the basic BAR level.

The 2016 zoning reform had the goals of providing dignified housing, guiding urban growth, improving urban mobility, improving life in the neighborhoods, promoting economic development, incorperating an environmental agenda, and preserving cultural heritage. A key feature of this reform was to standardize building parameters across the whole city by assigning a fixed set of building parameters to each zone type. Under the new reform, blocks were assigned to a zone type, and then the city-wide building parameters associated with that zone type would be applied consistently across the city. This reform removed the power of local subprefeitura governments to set BAR and other building parameters within their jurisdictions. In addition to that, it simplified and centralized the BAR regulation, by implementing a Basic BAR of 1.0 in the entire city. The 2016 reform also eliminated the district-level maximum amounts that could be built above the basic BAR level.

A main idea in the 2016 law was to group zone types in to three major groups corresponding to a particular development strategy. Every zone type was labeled as one of the following categories: transformation, qualification, or preservation. The goal in transformation zones was to promote higher urban density, in terms of both residential and non-residential structures, near the city's main transportation corridors. The aim was to reduce the city's traffic and bring people closer to their jobs by improving land use in areas closer to medium and high public transportation networks, such as train, subway, monorail, and bus corridors. The objective in qualification zones was to improve life in residential neighborhoods by favoring moderate urban density; the standard max BAR for these zones was set at 2.0. Zones were designated "preservation" status with the purpose of preserving the environment and cultural heritage of the city. Figure A1 shows

[7] The law also required developers to pay extra fees to build above the newly established basic BAR levels. A building developed with a BAR between the basic and max BAR was required to pay the so called "onerous grant" fee, charged per square meter built above the basic BAR level. Buildings developed between the minimum and basic BAR did not have to pay any extra fees.

a color map of the 2016 reform, with the transformation zones in maroon, qualification zones in gray and yellow, and preservation zones in green. Figure A2 displays the striking detail of the block-by-block land use regulation in the middle class neighborhood of Jabaquara. The 2016 zoning reform is valid for 16 years. Given that much of the predicted worldwide urban growth will occur in developing countries, it is interesting to test whether a zoning reform of this type, in a developing country city, actually leads to changes in the built environment and improvements in resident welfare.[8]

3 Data

3.1 Zoning data

Zoning data comes from the Cidade de São Paulo Desenvolvimento Urbano. We geo-reference and digitize maps of zoning boundaries at the block level for 2004 and 2016; there are 22 zone types in 2004, and 38 in 2016. In total, we are able to match 45,082 of São Paulo's city blocks to a zone-type, or 96% of the city's 46,987 blocks. We then match these zone-types to the relevant minimum and maximum allowable density parameters at the neighborhood and zone level. From this, we calculate the maximum allowable BAR. In some cases, these parameters vary within a zone-type depending on the size of the lot. In these cases, we define the maximum allowable density parameter as the maximum of all possible values.[9]

We obtain the max BAR values in both periods for 43,250 city blocks.[10] The underlying BAR variation for each block is mapped in Figure 1, along with the borders of the city's subprefeituras. More than 50% of all blocks experienced an increase in max BAR. There is a somewhat general pattern of blocks in the outskirts of the city experiencing positive changes in maximum allowable BAR, and blocks within the central regimes experiencing negative changes or no change in maximum allowable BAR. However, there are many blocks in the central area that received positive (green) changes in their maximum allowable BAR - typically along major transportation corridors.

[8] The reform may be toothless if, for example, developers can evade zoning rules via paying bribes – in such an environment zoning rule changes would have little impact as even existing rules are not enforced. Also, it is possible that zoning rule changes translate in to only minor built environment changes because of other market frictions, such as problems in land acquisition (Bryan et al., 2017) or developer credit constraints.

[9] In defining the max BAR before 2016, we account for the fact that some blocks are in districts where the allowed capacity for building above the basic BAR level has been exhausted. In these blocks, the "effective" BAR prior to the 2016 reform is the basic BAR level.

[10] Note that this is less than the 45,082 city blocks for which zoning information is available. These missing blocks are primarily parks, municipal areas, and bodies of water. The remainder are cases in which zoning information was available, but BAR parameters were missing or not relevant for the particular category.

Figure A3 presents binned scatter plots on the relationship between block level characteristics (x-axis) and the change in maximum allowed BAR from the 2004 to 2016 zoning regime. On average, blocks with lower levels of BAR in the 2004 regime received greater increases in allowable BAR. We also see that blocks with higher residential shares of constructed area received greater BAR increases. Blocks with greater density (as measured by the log constructed area per square kilometer) received on average lower changes in BAR, and blocks with greater average land values also received lower changes in BAR.[11]

3.2 Permitting data

Data on building permits were scraped from the Urbanism and Licensing Center of the City of São Paulo, or SMUL. Building Permit data is useful in tracking where and at what volume development is taking place in the city of São Paulo. There are over 50 different types of permits issued through SMUL, such as permits for demolition, installation of security systems, etc. but the majority of them are related to the construction of new buildings. The data is organized by quarter and by year of its filing and includes information about which region of the city the permit is in, the address, the zoning and land use classification, as well as the engineers, architects, and owners leading the project.

We obtain data on approximately 30 thousand total residential permits for the period 1997-2020, of which approximately 85% are multifamily buildings and 15% are single family dwellings.[12] We then aggregate the number of permits and the number of units at the quarter-block level, yielding a panel of 3,195,116 quarter-block observations. The main outcome variables are the quarterly count of total new building permits for single and multifamily buildings, plotted in Figure A5. Permitting activity was much larger in the mid-2000s, and declined since 2014 due to the national economic and political crisis. Interestingly, multifamily permits filings more than doubled since the approval of the 2016 reform.[13]

[11] As with most major zoning reforms around the world, the Sao Paulo 2016 reform changed multiple building parameters across space. In Appendix A we discuss other zoning parameters and our decision to focus on BAR changes within this reform.

[12] The sample size of commercial permits is not large enough to support our research design, so we do not analyze them in this paper. See Figure A4 for the count of commercial permits over time.

[13] We note that Sao Paulo does have an informal construction sector and that there could be interactions between the development we study and building in the informal sector. Given a lack of administrative data on informal developments and prices we leave an analysis of the informal sector response to the zoning reform for future work.

3.3 IPTU data

Information on the stock of buildings in São Paulo, including constructed area, lot sizes, assessed construction value, number of units per building, and assessed land value comes from the IPTU property tax data, which we obtain annually from 1995 to 2019. In 2016, the year of the zoning reform, this data covers 3,316,608 individual tax paying units in 1,582,532 unique buildings. The construction and land values are assessed values produced by the São Paulo property tax assessor office, and form the basis for annual property tax payments; these values are only indirectly based on market transactions. We collapse this data to the block-level to obtain block-level average lot area and constructed area in m^2, as well as mean land value and construction value per m^2. We also obtain the share of lot and constructed area with residential vs. commercial designation by the property tax authority, as well as the total number of units and buildings in each block. This information is available in the pre-2016 period for a total of 43,990 São Paulo city blocks, or 94%.

3.4 Listings and Price Data

Given that the IPTU value data is based on assessments as opposed to market values, we also collect listing price data from two different sources. The first source is the online marketplace Properati. The data has unique entries for approximately 200,000 buildings listed for sale and for rent in 2016. Each listing in the system contains the price, the type of transaction (rent or sale), the location, the type of unit (i.e. apartment, house, office, etc.), and a general description of usage (residential, commercial, bare-land or non-specified). The IPH (Hiperdados-Properati Index), which used data from Properati, was the most complete real estate pricing indicator in Brazil until it was sold and renamed to Casafy (who stopped publishing the data in 2016).

DataZap, a real estate analytics company that also has public online interfaces similar to Zillow, provided prices and listing for more recent periods. From DataZap we obtained data containing aggregated block level information on the number of listings, average price of listings per block, and average characteristics of listings (such as number of rooms, square meters) for the years 2019-2022.[14] The data separately covers apartments and single family homes, and within these categories data is provided for sales and rental listings.

[14] DataZap cannot share its microdata because of confidentiality issues. Since DataZap counts listings on a monthly basis, from three different sites, we correct for this overestimate by dividing the total number of listings in each block by 15, which is the average number of months listed (6) times the average number of sites listed (2.5)

3.5 RAIS Data

We measure labor market economic activity at the block level using the total formal sector wages paid to workers whose firm address is within a block. We obtain this variable from the RAIS data, which is individual level monthly wages data for all formal sector workers in São Paulo. We aggregate the individual monthly wage data up by year, and then further aggregate at the block level to obtain total annual wages paid per block. This data is available from 2000-2019.

3.6 Commuting Zone Survey Data

We use commuting survey data from the "Pesquisa Origem e Destino 2017" survey. The commuting survey covered São Paulo metropolitan statistical area households and includes information on demographics, place of residence and place of work. We take as our sample the 24,800 households for which the household head is working. The survey was stratified by commuting zones, and for this employed sub-sample we obtain 492 commuting zones within the São Paulo metropolitan area (i.e. including both the São Paulo municipality that we study as well as the surrounding suburbs). 329 of these zones are within the municipality itself. From this data, we take individual home and work locations, which we use to calculate commuting distances, household head education and age, household size and total monthly income, and dwelling ownership status. We also use this data to estimate several commuting zone-level characteristics, including average income, education, and the share of paved roads.

4 Border Discontinuity Design and Results

4.1 Empirical strategy

Our primary identification strategy to estimate the impact of zoning reform is to compare the evolution of permitting activity in geographically close blocks with different treatment status. We define "treated" blocks as blocks in zones that experienced an increase in maximum allowable BAR as a result of the 2016 reform. Control blocks are those that fall in zones which experienced either no change or a reduction in max BAR. Then, for each treatment block, we calculate the distance in kilometers to the nearest control block. Distance is calculated from the block's centroid to the edge of the nearest block in the other group. This distance defines the running variable in our border discontinuity design for treatment blocks. The running variable for control blocks is the (negative) distance to the nearest treatment block. Our boundary discontinuity design focuses on outcome comparisons between control blocks with small absolute values of the running variable (i.e. control blocks near to treatment blocks) and treatment blocks with small values of the running variable (i.e. treatment blocks near control blocks). In the Appendix we report the number of blocks in 50 meter bins of the running variable, and find no discrete jump around the cut-off (Figure A6.) If the zoning reform had a causal effect on building activity we would expect an increase in permitting activity at the boundary after the implementation of the reform; by conducting the boundary discontinuity design before the reform, we can assess whether treatment/control blocks were on similar permitting trends prior to the reform.

The core treatment variation is the change in maximum allowable BAR in nearby treatment and control blocks. Figure 2 shows how the change in BAR from the 2004 to 2016 zoning regime varies as we move towards the 2016 zoning borders from control blocks to treatment blocks. The figure shows averages of max BAR by bins of .1 km distance to the closest border, i.e., the closest block from opposite treatment status. City blocks just to the left of the cut-off experienced an approximate .1 decrease in their maximum allowable BAR, while blocks just to the right of the cut-off experienced an approximate 1.3 increase in their maximum allowable BAR. In appendix Figure A7 we show that the first stage is almost identical when restricting the control group to blocks with no change in max BAR. Figure A8 reports the average max BAR values in our treatment and control blocks before and after the 2016 reform, not just the change in max BAR shown in Figure 2. Treatment blocks had lower max BAR values prior to the reform relative to control blocks, and have higher max BAR values after the reform. The pre-existing differences prior to the reform are to some extent mechanical, in that treatment blocks are defined as those which experienced an increase in BAR in the 2016 reform. But the fact that max BAR did differ prior to the reform, even in a narrow bandwidth around the border, strongly suggests that we should focus on how outcomes change before and after the reform, as opposed to just analyzing a cross-sectional border discontinuity design after the reform.[15]

[15] Aside from the pre-reform max BAR difference across the treatment/control boundary, we do not find major differences in other building characteristics across the boundary. Figures A9 and A10 assess the covariate balance of existing structures and economic activity across treatment and control blocks in the year prior to the zoning reform. The x-axis groups blocks within .1 kilometer bins away from the cut-off, and the y-axis plots the binned-average change in covariates. Regarding the built environment, we find no meaningful differences in block density (constructed area per square meter), average land or constructed value, number of buildings, or residential/commercial share of constructed area at cut-off. We also find no meaningful differences in total employees in private firms, mean per-worker wages, total wages paid, and the number of private firms. Table A1 presents regression discontinuity estimates of these differences. Out of twelve comparisons we find one variable (log of aggregate private employee wages) is statistically significant at the 10% level or higher. Based on Figure A10 this statistical difference does not correspond to a jump at the cut-off that

To estimate the treatment effect of higher BAR levels, we use the following regression discontinuity model:

$$y_{ij} = \beta 1\{x_{ij} > 0\} + f(x_{ij}) + \delta_j + \epsilon_{ij} \tag{1}$$

where y_{ij} is an outcome in block i which is located in subprefeitura j, x_{ij} is the value of the running variable for block i in subprefeitura j, δ_j is a subprefeitura fixed effect, and ϵ_{ij} is an error term. The indicator function $1\{x_{ij} > 0\}$, designating treated blocks, is our main independent variable of interest, and β is our estimate of the treatment effect. The function $f(\cdot)$ is a polynomial of distances fully interacted with $1\{x_{ij} > 0\}$. We cluster standard errors at the commuting zone level; there are 329 commuting zones in our data so this clustering is substantially more conservative than clustering at the block level.[16] Commuting zones are also a natural level to cluster given this will be our neighborhood unit of analysis in the supply/demand model that estimates welfare effects. In the main results we use all of the data to estimate the regression discontinuity model (i.e. we do not limit the bandwidth); later we show robustness to smaller bandwidths.[17]

Table 1 reports RD point estimates and standard errors for our "first stage" using four versions of this specification with change in max BAR as the outcome variable. Column 1 compares all treatment versus control blocks (i.e. only includes the indicator $1\{x_{ij} > 0\}$ in the model), column 2 adds a linear control of the running variable (and interacts it with a treatment indicator), column 3 adds quadratic controls, and column 4 adds cubic controls. RD estimates show a stable 1.4 point estimate, which is large relative to the mean change in max BAR of -.153 in the control blocks within .1 km of the boundary. In Panel B of table 1 we add subprefeitura fixed effects, so that all variation comes from changes within subprefeitura. Estimates remain practically unchanged, suggesting that the reform treatment is not driven solely by pre-reform differences in how different neighborhoods controlled zoning parameters.

is economically meaningful relative to the bin by bin variation.

[16] Our main estimates cluster standard errors at the commuting zone level. To check the robustness of these results to spatial correlation in errors we also estimated Conley (1999) standard errors, allowing for spatial correlation within 1, 5, 10 and 20 km radii around each observation. Overall the standard errors get larger as allow for correlation of errors at larger distances, but the statistical significance of the main estimates remains.

[17] We present our main results using the full bandwidth because of potential measurement error in calculating distances. Measurement error can occur for two reasons. First, there are very few small blocks with inter-block distances that are less than .05km (50 meters), as seen in the histogram Figure A6. Second, while the 2016 map has quite precise geo-location of blocks, the 2004 reform was digitized from historical pdf files which can lead to errors. Since treatment assignment is based on changes in Max BAR, treatment assignment at small distances are more likely to suffer from mis-assignment of blocks to treatment and control groups.

4.2 Reform Short-Run Effect on Building Permits

We now use the same set of treatment and control blocks to estimate the causal effect of the greater allowable BAR on building permits. Figure 3 splits our permit outcome variable into multifamily permits (top two figures) and single family permits (bottom two) figures. The outcomes represent average quarterly building permits in a block, and the left panels show pre-reform data (2012q2 - 2016q1). Focusing on multi-family permits first, the pre-reform panel shows a small difference in permits at the discontinuity. In contrast, the post-reform period (2016q2 - 2019q4) show that treatment blocks just on the higher BAR side of the zoning border experienced approximately 0.004 more permits issued relative to the control side of the border (0.0085 versus 0.0045). The higher BAR allowance causes almost a doubling of average multi-family permits filed per quarter relative to control blocks.[18]

Interestingly, the zoning treatment effect on permits is concentrated in multi-family units. The bottom two panels of figure 3 show that not only are total number of single-family permits smaller in both pre- and post-reform periods, but also that there are no differences around the spatial discontinuity. Overall, the causal effects of the reform are concentrated in multifamily permits, consistent with these buildings being more sensitive to BAR constraints.[19]

Our key identification assumption is that, in the absence of the 2016 zoning reform, outcomes would have evolved similarly across the zoning borders at which the change in max BAR switches from positive to non-positive. Figure 4 presents separate regression discontinuity estimates for each half-year, including both pre- and post-reform quarters. The RD estimate is generally not statistically significant prior to the reform. We see the treatment blocks experiencing greater permitting activity approximately two half-years after the reform, and the point estimates more than double three years after the reform. The emergence of the treatment control differences approximately one year after the reform is consistent with a causal effect of the reform, as opposed to pre-existing differences correlated with zoning changes.

[18] Figure A11 tests if the multi-family permits point estimates, pre- and post-reform, change according to how much of the sample we use away from the cut-off (i.e. considering larger bandwidths). Post-reform estimates with fixed effects are consistently around .002 and .003, independent of the bandwidth size. We also have estimated these models using the optimal bandwidth procedure in Calonico et al. (2017). Across the specifications the optimal bandwidth is quite small at .0954 km (95 meters). We find very similar point estimates and statistical significance levels using uniform kernels conventional, bias-corrected, and robust inference procedures. The effect sizes decline to .00223 extra permits per quarter and are significant at the 10% level using a triangular kernel; this may be because the triangular kernel weights observations near the cut-off strongly, and these observations may have significant measurement error in their treatment control classification.

[19] Both pre- and post-reform figures show a downward slope in permitting activity as we move from deeper in the control area towards the cut-off, and then a slightly positive slope from the cut-off towards deeper in to the treatment area. We find a similar pattern in the total number of buildings per block in the 2015 IPTU data (see Figure A9).

Table 2 reports RD estimates for multi-family permits, in analogous form to the presentation in Table 1. Panel A reports RD estimates ranging from 0.0044 to 0.0048 (excluding the Column (1) estimate which does not control for the running variable). The inclusion of subprefeitura fixed effects in Panel B has the impact of reducing the magnitude of causal estimates to a range of 0.0023 to 0.0031. This means that part of the permitting effects are explained by differences in subprefeitura boundaries.[20] Combining the reduced form estimates with the first stage max BAR treatment magnitudes estimated in table 1 reveals that increasing max BAR by 1 leads to an increase in multi-family permits between 28% - 63%, relative to nearby control blocks. [21]

Figure A12 presents regression discontinuity estimates on the treatment effect of allowing greater BAR ratios separately for blocks with below and above median land values. The binned averages in these plots are produced by first splitting the sample in to below and above median groups, and then using the distance to the nearest zoning border as the running variable. The figure suggests that the zoning treatment effects are largest in areas with higher pre-existing land values. In a regression discontinuity where we interact the treatment variable with a dummy for the block being above the median land value we find that the increase in permits is .00121 permits per quarter higher for treatment blocks with above median land values, which is 50% higher than the effect for treatment blocks below median land value. This difference, however, is not significant at the 5 percent level – although a linear specification does find significant differences.

Our short-run analysis has focused on changes in BAR in nearby blocks; given that the reform included both changes in BAR and zone type designations by block, it is useful to characterize how nearby treatment and control blocks differ in zone type designations as well. Appendix Figure A13 and Appendix Table A3 show that there are *no* differences in the share of blocks that are designated for residential use - this could potentially confound our results if more treatment blocks had changed from industrial to residential use, for example. But as the law intended, we do find small increases in the share of treatment blocks designated as transformation zones, while control blocks have small increases in the share of blocks designated as qualification and preservation.[22]

[20] Appendix Table A2 applies the same RD strategy in a Poisson estimation, finding similar results.

[21] For brevity we do not include fuzzy RD estimates because our counterfactual results ultimately depend on our full supply model estimated later – where we present estimates on the instrumented BAR variable.

[22] We also looked at other zoning parameters that could be changing alongside max BAR. Appendix Figure A14 and Appendix Table A4 show RD results for building height, basic BAR, shadow ratio, and maximum lot size. Again, as expected, we find small discontinuities in all variables. More interestingly, we re-estimate our main RD specification for both new multifamily units and max BAR change, now including those additional zoning parameters as additional controls. Appendix Table A5 shows that all treatment coefficients drop by about 1/3. However, when dividing the new

4.3 Spillover Effects?

We identify the short-run treatment effect of relaxing zoning rules by comparing areas that received a higher max BAR designation to nearby areas that did not receive such treatment. In this section we analyze the extent to which increases in allowable BAR levels might affect nearby blocks (i.e. spillover effects). On the one hand, developers may act independently, and spillover effects might be small during the post-reform period when developers are simply filing for permits. But there are at least three reasons to estimate spillover effects in our context. First, projects could move from nearby control blocks to treatment blocks, leading us to over-estimate the effect of the BAR reform. Second, buildings in nearby control blocks could *increase* if the greater expected density in the nearby treatment blocks make nearby control blocks more attractive areas to develop as well. Finally, part of the treatment effect in treatment blocks could reflect agglomeration benefits of other nearby treatment blocks - although these benefits would likely appear with long lags.

We follow a simplified variant of the methodology in Diamond and McQuade (2019) to assess the importance of spillovers in the context of a highly localized zoning reform. The basic strategy is to treat both treatment and control blocks near a zoning boundary as “treated,” in the sense that they are nearby to an area where a major zoning change occurred. We compare outcomes for these blocks near boundaries to a “pure control” set of blocks that are farther (i.e., greater than half kilometer) away from the boundary, before and after the reform. To operationalise this spatial difference-in-differences strategy we estimate the following regression model:

$$y_{it} = \sum_{j=1}^{5} I(t > 2016Q2) * I(dc_{j*-.1} = 1) + \sum_{j=1}^{6} I(t > 2016Q2) * I(dt_{j*.1} = 1) + b_i + q_t + \epsilon_{ij} \qquad (2)$$

where y_{it} is the number of permits issued in block i in a quarter t, $I(t > 2016Q2)$ is an indicator for post-reform, $I(dt_{j.1} = 1)$ is an indicator for a treatment block that is $j * .1$ km away from the nearest control block, and $I(dc_{j*-.1} = 1)$ is an indicator for a control block that is $j * .1$ km away from the nearest treatment block. The equation also includes block and quarter fixed effects. The omitted category consists of control blocks more than .5 km away from the boundary, and all treatment blocks with distance greater than .5 km are bunched in a 0.6 bin.

reduced form effects by the new smaller first stage, we find results similar to the baseline estimates shown in Tables 1 and 2. In the complete supply model described in the next section we will estimate development activity as a function of all local zoning regulations, not just max BAR.

Figure A15 presents the estimated coefficients. On average permits are higher in all treatment blocks relative to all control blocks. The difference between treatment and control groups is slightly larger within .1 km to the boundary. But they are not statistically different from each other, in part because the distance bins are quite small. Moreover, we do not observe control blocks near treatment blocks appearing to have particularly low permit averages. Relative to the omitted category, the distance of a given treatment (control) block away from the boundary does not appear to be strongly associated with permitting activity. The finding of non-existent spillover effects in building activity is consistent with Turner, Haughwout and Van Der Klaauw (2014) finding of small effects of land-use regulations on nearby property prices.

4.4 Permit Approvals and Future Construction

In addition to permit filings, the city of Sao Paulo also provides data on approved permits. Following a similar spatial regression discontinuity, we look at differences in the number of building permits that were actually approved by the city. It turns out that more than 75 percent of permits are approved within 3 to 4 years of their filings, a number that only slightly increased after the 2016 reform. In Table A6 we estimate a similar RD for total multi-family permits that were approved after the reform, and find estimates in the range of 47-88 percent increase in approved permits for each 1 unit of max BAR increase. Those ranges match our estimates for filed permits.

Given the lags involved in filing and approval of permits, and the additional years necessary to observe any type of multifamily projects - which generally have extra construction delays - using constructed buildings as an outcome would give a biased picture of the zoning reform's full effects, at least in the medium run. However, we can use data prior to the reform to estimate the conversion rate of permits to constructed buildings, in order to help quantify how our results on permits are likely to convert to constructed buildings. Figure A16 estimates a block-level event-study model on the impact of a permit being issued on the density of new construction, measured as new constructed area divided by total land area in the IPTU data, in the period prior to the reform (2004-2016). The figure suggests that a new permit issued in a block is correlated with increases in density up to 15 years after the issuance. Of course, these calculations are only useful to the extent that past relationships between permits and finished construction hold. We will use this relationship to translate the completely observed short-run permit effects into construction effects when we estimate our structural model to evaluate welfare.

4.5 Reform Medium-Run Effects on Availability of Housing and Listing Prices

So far we have analyzed the reaction of developers to zoning reform based on the fastest outcomes we can observe, i.e., the filings and approvals of zoning permits. Now, we look at medium-run effects of the reform on different outcomes.

First, we use DataZap information to test if the availability of housing units for sale actually increased by 2022, six years after the reform. Due to only having listing data from 2019 onwards, the best we can do is estimate whether the growth in listings from 2019 to 2022 at the block level was faster in treated versus control blocks. Given the large lags involved in applying for a permit, getting an approval, and then constructing a property, we argue that year 2019 is actually a reasonable baseline to observe the potential effect of the reform on housing availability. But we caveat this result by noting that we cannot test for parallel trends in this analysis, because we do not have pre-reform data on listings. Figure 5 shows the boundary discontinuity estimates for the change in listings of properties for sale from 2019 to 2022, while Table A7 provides point estimates under different specifications. In our preferred model we find the max BAR treatment effect equal to an increase of 2.1 additional listings from 2019 to 2022, which is a 10.3% percent increase relative to the average number of listings in 2019 in the bin of control blocks just to the left of the cut-off. We find relatively similar results with difference bandwidth choices.[23]

We do not apply a similar border discontinuity design for listing prices. Housing markets are generally not segmented by blocks, and therefore availability of homes for sale in block may impact prices in a broader neighborhood area - prices in fact may evolve smoothly around zoning boundary discontinuities. Instead, we aggregate our listings by the 329 commuting zones and compare number of listings and listing prices by that geography. Number of listings is just aggregated by commuting zone, but prices are aggregated in three steps: 1) Calculate residual prices at the block level by regressing listing prices of a given type of property (apartments or single-family) on listing characteristics for that type of property; 2) Calculate average residual price at block level using weights based on number of listings; 3) Aggregate block level residual prices at the commuting zone using weights based on the stock of residential properties in each block

[23] In Table A8 we implement the optimal bandwidths as suggested in Calonico et al. (2017). The optimal chosen bandwidth here is very small at 55 meters, which may increase measurement error as discussed above. The point estimates are positive and economically meaningful (equal to an approximate 5% increase in permits relative to the level of permits in the control bin closest to the cut-off), but are not statistically significant at the 5% level. Figure A17 presents treatment effects and confidence bands for all bandwidths above 100 meters. Overall the results are consistently positive for small bandwidths (i.e. they do not collapse to zero at the small bandwidth of 100 meters), and get somewhat larger as we increase the bandwidth beyond 300 meters.

based on the official property tax data. Max BAR is also aggregated at the commuting zone level by averaging the Max BAR changes at the block level.

Figure 6, left panel, shows that commuting zones with larger average increases in max BAR also had larger growth in availability of homes for sale. We formally test this relationship in a difference-in-difference regression of changes in outcomes on changes in max BAR, which holds constant fixed features of commuting zones. We find that a 1 point increase in max BAR is correlated with a 10.9 percent increase in listings growth, which is remarkably similar to the causal estimates for that variable using the boundary design at the block level. The right panel shows a similar comparison for listing prices. Now, the slope is negative, with a 1 point increase in max BAR being correlated with a 5.7 percentage point reduction in home prices. Estimates for these zone-level regressions are presented in Table A9.

In summary, these short and medium-run outcomes suggest that the reform had strong effects on blocks and neighborhoods that were allowed to have more densification. Developers immediately applied for more new permits in those blocks, and also received more new approvals a few years later. Subsequently, those constructed units were listed for sale in blocks with higher max BAR, increasing housing supply. When aggregating listings at the neighborhood level, the increased supply was accompanied by reductions in prices. Those reduced form results suggest important changes in welfare for city dwellers. But since other aspects of neighborhoods may be changing, such as the built environment and congestion, it is important to include all these factors in a model to fully understand the magnitude of the welfare effects and its distributional consequences.

5 Welfare Evaluation

In this section we estimate an equilibrium model of supply and demand to predict the long-run welfare consequences of the Sao Paulo 2016 reform. Our supply model is an extension of the reduced form results presented above, and allows for endogenous zoning regulation and prices. On the demand side we follow the standard neighborhood choice approach, where individuals maximize utility when choosing neighborhoods as a function of observed and unobserved location features. The richness of the Sao Paulo data allows us to estimate the value of both positive and negative aspects of more densification in each neighborhood. We then back out equilibrium prices under the assumption that quantity supplied equals to quantity demanded at the neighborhood level, and simulate welfare and distributional consequences under alternative zoning scenarios.

5.1 Model of Residential Supply

We model the construction of new residential housing as an exponential function of housing prices, building density restrictions, and other location characteristics. The Poisson functional form is used because of the strictly positive, discrete count nature of housing units. In particular, we estimate the following supply equation for location j:

$$E[s_j | p_j, M_j, X_j^s] = \exp(\alpha^s p_j + \psi M_j + \beta^s X_j^s) \tag{3}$$

where s_j is the total number of building permits in location j, p_j is the average residential 2016 listing price in j, M_j is the maximum allowable BAR in neighborhood j, and X_j^s is a vector of other housing and regulatory characteristics of location j that affect development activity. These other variables include construction density in j, the average building age, the average number of units per building, and the average value of the pre-2016 BAR zoning parameters, and the 2016 non-BAR zoning parameters (maximum shadow ratio, minimum and basic BAR, maximum height, minimum and maximum front setback, and maximum area), averaged across all blocks in j.

We estimate the supply model at the subprefeitura-quantile-level, where city blocks are aggregated into 40 quantiles of our regression discontinuity running variable (i.e. distance to the 2016 zoning change boundary) within each subprefeitura. This aggregation is done for two reasons: First, price information is relatively sparse at the block level, and so aggregating allows us to get price information for all the observations. Second, the distance quantiles allow us to still exploit the boundary approach to identify the Max BAR effect. The estimation procedure will focus on comparisons between permit outcomes for the quantiles just above and below the BAR increase cut-off within subprefeituras. Summary statistics for our supply variables are shown in Table A10.

5.1.1 Instrumental variables

In equation 3 we treat p_j and M_j as endogenous neighborhood characteristics, while assuming other physical location characteristics are exogenous. The assumption is that developers take these other characteristics as given when making development decisions, and do not adjust their estimate of these characteristics when making supply decisions over the 10 year horizon. This assumption is reasonable because changes induced by new construction to these variables are relatively small, given these variables are generally characteristics of the total stock of buildings. We caveat that the plausibility of this assumption depends on the size of the reform we are evaluating; while credible for the 2016 change in policy, other larger reforms could affect stock variables more significantly. To address endogeneity in the price p_j and max BAR M_j, we estimate the supply model with multiple instruments using the GMM estimator of Mullahy (1997) with additive errors to form moment conditions. The instrument set is $W_j^s = [T_j, B_j, X_j^s]$. The first excluded instrument T_j – used for max BAR – is an indicator for whether the subprefeitura-quantile j is treated by the 2016 reform. To leverage the regression discontinuity design for identification of ψ, we also include distance to the boundary D_j, and $D_j \times T_j$ as control variables in X_j^s. Validity of the max BAR instruments was discussed in detail in the previous sections.

The second excluded instrument, B_j, instruments for market price. A major challenge in the housing supply literature is the identification of supply elasticities with respect to price. Inspired by Baum-Snow and Han (2024), we instrument for price using a Bartik-style demand shock. The idea is to exploit plausibly exogenous variation in exposure to national labor demand growth for neighborhoods based on the historical presence of sectors that experienced substantial national growth. Formally, let B_l be our Bartik instrument defined as:

$$B_l = \sum_k z_{lk} g_k^{-SaoPaulo}.$$

where k indexes one of the 59 economic sectors from the RAIS data, and l indexes commuting zone. $g_k^{-SaoPaulo}$ is the national growth rate in employment, excluding the Sao Paulo municipality, in sector k from 2007 to 2017. z_{lk} is the commuting zone level sector k share of employment in 2007 (the "initial share"); z_{lk} is calculated as $\frac{L_{lk}}{\sum_k L_{lk'}}$, where L_{lk} is the total number of formal sector employees in the RAIS data in commuting zone l in sector k, divided by the sum of employees in that commuting zone excluding the sector. [24] Lastly, we aggregate the commuting-zone-level Bartik shock, B_l, into a subprefeitura-by-quantile shock using the weights in Appendix B.

Our identification assumption is that our Bartik instrument only affects new permit issuances in 2017 through its effect on prices in 2017. The threat to identification is that the initial shares in 2007 could be related to 2017 permit issuance levels through some non-price mechanism conditional on our control variables. Table A11 presents correlations between subprefeitura by quintile

[24] Appendix figure A18 shows how these initial shares are distributed across our 329 commuting zones. Retail tends to have the largest share of employment in 2007. Appendix Figure A19 shows national employment growth rates from 2007 through 2017 for the 59 economic sectors.

region labor shares in the top 5 most important sectors in our Bartik instrument with neighborhood characteristics. We note that the R^2 statistics here are above .24 for three of five sectors, i.e. sectoral shares are not randomly assigned in our setting so it is possible that the exclusion restriction of this instrument fails. Nonetheless, we follow Baum-Snow and Han (2024) in reporting this estimate as the latest available method for estimating housing supply elasticities.

Appendix Figure A20 shows a scatter plot of the first stage relationship between market prices in 2017 and our Bartik instrument. Each point is a commuting zone. There is a positive correlation; commuting zones that had higher initial shares in sectors that ended up growing faster nationally also do have higher price levels in 2017. The right panel of the figure shows a binned scatter plot version, indicating a strong first stage.[25]

5.1.2 Supply estimation results

We estimate the supply model using data on 5,375 new residential building permits filed with the São Paulo city government between 2016-2019. The right-hand-side variables are taken from either IPTU or the block-level zoning maps; all of these variables are averaged across blocks within the subprefeitura-quantile. In total, we obtain 1182 subprefeitura by quantile observations, of which 900 have any new residential construction permitting activity over this period.

Table 3 presents the estimated supply coefficients. Column (1) does not instrument for either endogenous variable, while column (2) estimates the model using only the regression discontinuity instrument for M_j. Column (3), our preferred specification, instruments for both M_j and p_j with the RD and Bartik IVs, respectively. The estimates in (3) imply that a one-unit increase in maximum allowable BAR at the mean leads to approximately 5.05 additional new building permits for the average unit.[26] This estimate is in line with the RD estimates from Table 2. The price coefficient in column (3) implies that a 1,000 reais increase in price (18.7% of the mean) is associated with 2.34 new permits on average.[27] Note that column (1) substantially underestimates both coefficients by ignoring endogeneity. Instrumenting for BAR with the RD nearly doubles the estimated coefficient. In addition, the price elasticity estimate rises from 0.149 to 0.415 from columns (2) to (3).

[25] Appendix Figure A21 shows the corresponding figures for number of new permits issued in the post-reform period against the Bartik IV (the "reduced form" relationship). Goldsmith-Pinkham, Sorkin and Swift (2020) recommend analyzing how the overall instrumented coefficient compares to the coefficients estimated based on each possible sector as an instrument. Figure A22 shows this plot for our Poisson model. We see that most sectors are close to the overall estimate, suggesting that the price elasticity coefficient is not determined by one particular outlying division.

[26] $5.05 = 4.55 * (\exp(0.747) - 1)$, where 4.55 is the average of the outcome variable.

[27] $2.34 = 4.55 * (\exp(0.415) - 1)$.

This suggests a downward simultaneous equations bias in estimating the supply elasticity.

In columns (4)-(5), we show that BAR effects are small and statistically insignificant for single family permits, but large and significant for multifamily permits. This is consistent with the reduced form results in Section 4. Both housing types, however, exhibit similar degrees of price responsiveness.[28]

5.2 Model of Residential Demand

5.2.1 Choice model

Our model of residential housing demand follows a standard discrete choice framework (Berry, Levinsohn and Pakes, 1995; Bayer, Ferreira and McMillan, 2007). Household i chooses between $j = 1, ..., J$ alternatives, where the alternatives are one of 329 commuting zones within the city of São Paulo. The outside option is living outside the city - in the suburbs - taken by roughly 40% of the individuals in our residential commuting data. Individual utility from choosing to live in zone j will be:

$$u_{ij} = \alpha_i^d p_j + \beta_i^d X_j^d + \gamma_i \tau_{ij} + \xi_j + \epsilon_{ij} \tag{4}$$

where p_j is the price of housing in zone j, measured as the average Properati listing price per square meter in commuting zone j. X_j^d is a K-dimensional vector of housing amenities of location j. This includes an index of residential commuter market access (RCMA) that measures the extent to which j is located near high-paying jobs, using travel times from a zone commuting matrix.[29] X_j^d also includes the average age, in years, of the stock of housing units, average number of units per building, the overall zone-level constructed area density – defined as the sum of all constructed area divided by the zone geographic area – the share of households with a paved road, the average

[28] In the Appendix, we consider several robustness tests for the supply model. Our additive error structure, following a standard nonlinear regression setup, implies as an identifying assumption that the instruments W_j^s are exogenous to the exponential function of the endogenous variables. Table A12 allows the error term to enter multiplicatively instead of additively in forming the GMM moment conditions. The results are similar. Table A13 considers approvals, rather than all permit filings, and finds nearly identical magnitudes. Finally, Table A14 estimates the model with a linear 2SLS specification, allowing us to assess instrument strength with the first-stage Kleibergen-Paap F-statistic. The sign and significance of the coefficients is unchanged, and the instruments are highly correlated with endogenous variables, allaying concerns about weak instruments.

[29] We define the RCMA for zone j as:

$$RCMA_j = \sum_{i \in I_j} \frac{w_i}{d_{ij}}$$

where I_j is the set of zones that have at least one worker living in j and w_i is the average wage in i. Iceberg commute costs are modeled as $d_{ij} = \exp(\kappa \tau_{ij})$, where τ_{ij} is the average reported travel time in minutes between i and j. We set $\kappa = 0.01$, following micro-estimates from Tsivanidis (2022).

zone income, and the share of zone adults with a college degree. The term τ_{ij} measures commuting costs as the predicted travel time in minutes between the zone in which i works, taken as given, and zone j.[30] Finally, ξ_j is an unobserved location-specific "structural" utility shock, which may be correlated with price - we assume the variables in X_j^d are exogenous. The shocks ϵ_{ij} are distributed i.i.d. type I extreme value. The utility of the outside option is normalized to zero, $u_{i0} = 0$.

Let Z_i be a C-dimensional vector of household characteristics, including household size, age of the household head, a rental indicator, household income, and a college indicator for the household head. The heterogeneous demand parameters α_i^d and β_i^d take the form:

$$\begin{pmatrix} \alpha_i^d \\ \beta_i^d \\ \gamma_i \end{pmatrix} = \begin{pmatrix} \alpha^d \\ \beta^d \\ \gamma \end{pmatrix} + \Pi Z_i \tag{5}$$

where Z_i is the $C \times 1$ vector of demographic variables $z_{i,1}, ..., z_{i,C}$ and Π is a $(K+2) \times C$ matrix of coefficients containing: $i)$ $\pi_{\alpha,1} ... \pi_{\alpha,C}$, the interactions of price and each of the C demographics in Z_i $ii)$ $\pi_{\gamma,1} ... \pi_{\gamma,C}$, the interactions of travel time and each of the C demographics in Z_i, and $iii)$ $\pi_{\beta^k,1} ... \pi_{\beta^k,C}$ for $1, ..., K$, the interactions of all of the demographic variables with all of the K commuting zone characteristics.

We can then re-write utility separating out the parameters that only vary by neighborhood and the parameters that contain heterogeneity in preferences as follows:

$$u_{ij} = \delta_j + \mu_{ij} + \epsilon_{ij} \tag{6}$$

where $\delta_j = \alpha^d p_j + \beta^d X_j^d + \xi_j$ and $\mu_{ij} = (p_j, X_j^d, \tau_{ij}) \Pi Z_i + \gamma \tau_{ij}$.

Since we observe individual demographics, we can calculate the conditional choice probability that each individual i chooses option j. Define y_i as the choice indicator and θ^d as the vector of demand parameters. We have:

[30] We model the travel time between living zone j and working zone k with a regression of the form

$$\tau_{jk} = \mu + \eta \log(pop_j) + f(d_{jk}, \phi) + v_{jk}$$

where τ is the average reported travel time in minutes between the zones j and k, f is a polynomial function of distance d and regression parameters ϕ, and pop_j is the number of households in zone j. Population is included to allow for more dense areas to have longer travel times conditional on a distance to a given location. We estimate this equation on 30,934 route-level observations for which we observe any trip, using a cubic polynomial in distance. We then predict τ for all possible combinations of living zones in our model (329) and working zones reported in the data (517). This predicted value then enters the utility equation in the estimation based on individual i's working zone, taken as fixed.

$$Pr(y_i = j | Z_i, X, p, \xi, \theta^d) = \frac{\exp\left(\delta_j + (p_j, X_j^d, \tau_{ij})\Pi Z_i + \gamma\tau_{ij}\right)}{1 + \sum_{k \in J} \exp\left(\delta_k + (p_k, X_k^d, \tau_{ik})\Pi Z_i + \gamma\tau_{ik}\right)} \tag{7}$$

Note that the denominator of the conditional choice probability is taken over all locations in the city, implicitly assuming that all consumers choose from all possible neighborhoods in the city as well as the outside option.[31]

For estimation, we follow the standard two-step approach. In the first step, we use maximum likelihood to estimate heterogeneity parameters $\hat{\Pi}$, γ, and the fixed effects $\hat{\delta}_j$. The likelihood function is:

$$L(\Pi, \delta | X, p, Z, \xi) = \prod_i \prod_{j \in J_i} Pr(y_i = j | Z_i, X_j, p_j, \tau_{ij}, \xi, \Pi, \delta) \tag{8}$$

In the second step, we can recover the level coefficients by estimating $\hat{\delta}_j = \alpha^d p_j + \beta^d X_j^d + \xi_j$ using an instrumental variables regression. We folow Bayer, Ferreira and McMillan (2007) and create instruments for price by taking the average housing and spatial characteristics within a geographic "donut" around the neighborhood centroid. These IVs follow the logic of between-neighborhood competition: if location j is surrounded by higher quality zones, then j must lower its price to attract residents, implying a strong first stage. However, the nearest neighborhoods to j may create direct quality spillovers, violating the exclusion restriction. As such they are excluded from the calculation, creating a "donut" around j of neighborhoods that only influence the choice problem through their indirect impact on price in j.[32]

We follow Davis et al. (2021) in assuming exogeneity of neighborhood socio-demographics. While we allow households to have preferences over the pre-reform average demographic characteristic of neighborhoods, we do not allow for preferences over future demographic features of households who will move after the supply consequences of the reform. While this may be a restrictive assumption, below we show that the new extra housing units due to the reform represent a small fraction of the total housing stock, ruling out major demographic change. Moreover, as mentioned by Davis et al. (2021), this exogeneity assumption has two benefits: it avoids the search for additional instruments for demographics, and it resolves the potential problem of multiple

[31] In another version of the model, we restrict individual-specific choice sets choice set to the outside option and all of the locations in the consideration set J_i. The results are similar and available upon request.

[32] We select our IVs from a set of neighborhood characteristics that includes the paved road share, RCMA, housing stock age, average units per building, and density and spatial characteristics that include the favela share of zone area, flood-zone share of zone area, average slope, and metro station presence. We include neighborhoods from 5-20 miles from j in the average characteristics of competitors.

equilibria during the simulation exercise.

We estimate the model on a sample of 24,800 households with at least one employed member using the 2017 commuting survey data. To improve computational performance in the estimation routine, all variables are standardized by subtracting the mean value and dividing by the standard deviation across commuting zones. We estimate standard errors on all parameters by double-bootstrapping both the first and second stage of the estimation with 500 replications.

5.2.2 Demand estimation results

Table 4 presents our estimated demand parameters. The columns indicate the nine commuting zone characteristics (price, travel time, RCMA, age, units, density, paved, income, and education) - summary statistics of the main demand variables are in Table A15. We allow demand for each characteristic to vary by the following individual demographics: household size, age, renter, income, and college degree. Bootstrapped standard errors are in parentheses. All demographics and neighborhood characteristics are standardized, so base coefficients are interpreted as the preference for a one standard deviation increase in the neighborhood characteristic for the demographically "average" household. The base coefficients are taken from column (9) in Table A16, which compares several different specifications of the demand-side IVs.[33]

Focusing on these base coefficients, we find a negative elasticity of demand with respect to price and travel times. We find a positive elasticity with respect to RCMA, which implies that conditional on price and travel time to a given workplace, households prefer to live near areas with many high paying jobs. Regarding the built environment, we find consumers dislike old housing stock and also dislike greater density within a given building (units). However, they like denser neighborhoods (density) that have better infrastructure (paved). The base coefficient on income is negative and imprecise, suggesting there is an average preference for living in lower income neighborhoods - the majority of residents in Sao Paulo are lower income. The base coefficient on education is positive, suggesting a preference for living in higher education neighborhoods.

[33] Table A16 shows results from different instrumental variable specifications. Column (1) presents the OLS estimate, column (2) includes only the average X characteristics of competitors, column (3) includes only spatial characteristics, and column (4) includes both sets of IVs. Columns (5)-(9) contain different subsets of the most powerful IVs, as indicated in the Table footer. The price coefficient is smallest in the OLS regression, at 0.79, and increases in magnitude to roughly 1.5-3.9 for the IV the different specifications. This downward bias in the OLS estimate of the price elasticity is consistent with the standard simultaneous equations bias in supply and demand systems. The IV specifications vary substantially in strength, with the strongest being the parsimonious single-IV specifications using RCMA and density in columns (5) and (6). Still, the price coefficients remain relatively stable across IV models. Since using multiple IVs increases the amount of information used for identification – although at the cost of first-stage power – our preferred specification in (9) uses the subset of four jointly strongest instruments: the favela share, slope, RCMA, and age.

The other rows in Table 4 present the coefficient estimate on the interaction term between the neighborhood characteristic in the column and the household characteristic in the row. For example, the negative coefficient -.185 on College degree in the first column indicates that the price elasticity of college graduates is more negative than that of the average household. Some interesting interactions between neighborhood characteristics and demographics are as follows. Higher income and college educated households appear to be more sensitive to price changes in this sample, and have somewhat greater taste for local density. As far as taste for number of units in the building, older, richer, and more educated households have greater dis-utility from living in buildings with greater units.

5.3 Equilibrium

With the estimated supply side parameters $\hat{\theta}^s$ and demand side parameters $\hat{\theta}^d$ in hand, we can solve for equilibrium in the residential housing market. Our counterfactual exercises will consist of imposing an exogenous zoning map M (the densification policy experiment) and then solving the equilibrium to obtain $p(M)$, a J-vector of counterfactual prices $p = [p_1, ..., p_J]$ in each location j such that supply and demand are equated under M. We conduct the equilibrium analysis at the commuting zone level. We analyze the following map M scenarios: 1) A baseline equilibrium that takes observed zone-level market shares in 2016 (just prior to the 2016 reform), and estimated demand parameters and calculates the price vector necessary to equate supply and demand. 2) A counterfactual where we simulate the model for ten years given the 2004 zoning map. 3) A counterfactual where we simulate the model for ten years given the 2016 reform zoning map. 4) A "double BAR" counterfactual where we keep BAR at 2004 levels for blocks that in reality received lower BARs in the 2016 reform, and double the ultimate 2016 BAR for blocks that received an increase in BAR in the 2016 reform. In simulations (2)-(4) we run the supply model for 10 years to estimate long-run impacts of new housing supply on prices, residential sorting, and welfare.

To calculate the equilibrium prices for a given zoning map, we first take the estimated supply parameters $\hat{\theta}^s$ and use them to calculate $S_j(p; M, X^s, \hat{\theta}^s)$, the market share of total housing supply in location j for a given price vector, supply characteristics X^s, and zoning map M. Then, using the demand parameters $\hat{\theta}^d$, we calculate $D_j(p; X^d, \hat{\theta}^d)$, the predicted market share of location j given prices and the demand structure. The equilibrium condition is that supply equal demand in each commuting zone J:

$$S_j(p; M, X^s, \hat{\theta}^s) = D_j(p; X^d, \hat{\theta}^d) \; \forall j \in [1, ..., J] \tag{9}$$

We obtain a system of J-equations in J unknowns and search for the price vector p that solves the equilibrium system of nonlinear equations.

To calculate the commuting zone-level demand shares $D_j(p; X^d, \hat{\theta}^d)$, we must aggregate the individual conditional choice probabilities into commuting zone-level shares by integrating over the empirical distribution of demographics F_Z:

$$D_j(p, X^d, \hat{\theta}^d) = \int Pr(y_{it} = j | Z_i, X^d, p, \hat{\theta}^d) dF_Z \tag{10}$$

Calculating the zone-level supply shares $S_j(p; M, X^s, \hat{\theta}^s)$ is more complicated, since it requires solving two aggregation problems. First, our choice equation refers to new permits rather than the stock of buildings, so we must translate new building permits into market shares of total housing units. Second, our supply side equations are subprefeitura-quantile level and must be aggregated to the commuting zone level. For details on this aggregation procedure, see Appendix B.

Our equilibrium condition equates the *market shares* of each location as predicted by our estimated demand and supply models. Implicitly, this assumes that the population of the MSA will grow to meet the new housing stock built after a given shock to BAR. If this were not the case, then there would have to be real vacancies somewhere in the MSA after a positive housing supply shock in the model, since the total number of housing units would exceed the number of possible residents. In this sense, the equilibrium in shares is an "open city" model where new migrants are assumed to enter and fill in new vacancies. We believe that this is a more reasonable assumption for the city of Sao Paulo, which is the wealthiest city in Brazil and plays a somewhat similar role as that of New York City in the United States. Our method also assumes that the demographic characteristics of these new migrants are the same as existing Sao Paulo residents, which we find some empirical support in recent census data.[34]

One potential criticism is that this "equilibrium in shares" assumption may be restrictive and lead us to underestimate price effects. An alternative assumption on the other extreme is that the

[34] To get a sense of migration patterns to Sao Paulo we looked at the fraction of 2010 census respondents who lived outside of the municipality as of 2005. Only 3.56% of Sao Paulo city residents report having lived outside of Sao Paulo metropolitan area as of 2005 (average annual in-migration rate of .71%). Among those recent migrants, 17.6% have a college degree, which is approximately 4% higher than the average college degree holding rate. The average income of these migrants is R$900 reals per month, which is 10% less than the Sao Paulo city average. An even smaller share, .35% of Sao Paulo city residents, report living in the suburbs as of 2005. These migrants from the suburbs are 17 percentage points more likely to have gone to college and earn approximately R$600 more per month. To summarize, the 2005 to 2010 migration to Sao Paulo is small, mostly comes from outside of the metropolitan area, and the average characteristics are similar to the Sao Paulo city as a whole. Given those numbers, we also assume in our simulations that the new households migrating from new construction will reflect the demographics of the city as a whole.

São Paulo MSA is a closed economy, such that any counterfactual increases in housing supply in the city must be compensated by corresponding vacancies in the suburbs. This implies an equilibrium condition in *levels*, assuming no vacancies at $t = 0$. For each counterfactual, we calculate equilibria under both the shares and levels assumptions and interpret these as upper and lower bounds, respectively, on the true counterfactual prices.

In solving this model, we assume that all of the supply side dynamics are induced by the zoning change, and not by developers responding dynamically to each other's behavior nor the future path of prices. Similarly, while we allow for consumers to respond to changes in the built environment according to our demand model (i.e. the variables price, travel time, RCMA, age units, density and paved roads), we do not allow for consumers to be forward looking in how they choose neighborhoods.

5.4 Counterfactual Results

5.4.1 Equilibrium outcomes

We begin by calculating the implied prices that equate the observed commuting-zone market shares (supply) and estimated demand (based on our demand model) just prior to the 2016 reform (the "baseline scenario"). The purpose is to get a sense of how well the equilibrium model implied prices can replicate observed listing prices when using our equilibrium calculation procedure. Figure A23 shows the correlation between model-predicted prices for the baseline scenario and the observed listing price data from Properati. Our model prices do a good job of replicating the observed market prices, with an $R^2 = 0.75$. Figure A24 similarly plots the observed demographic composition of zones (log of average income and share of household heads with college education) in the data against the demographics that would be predicted by the individual-level choice probabilities of the model at baseline. The R^2 are 0.89 and 0.96 for log income and college-educated share, respectively.

Table 5 presents the main zone-level results on prices and quantities from our simulations. All results are based on simulated outcomes ten years after 2016. Column 1 presents equilibrium outcomes assuming the 2004 zoning stayed in place from 2016 to 2026, with an average max BAR of 1.55. Column 2 presents results for a simulation where zoning is changed according to the 2016 reform map in 2016, with a higher average max BAR of 2.09. Column 3, the "Double BAR scenario" keeps BAR at the 2004 level for those blocks that had a BAR decrease in the 2016 reform, and doubles post-2016 reform BAR for all blocks that received a BAR increase in the 2016 reform. Even under the Double BAR scenario, the average max BAR in the city is still just 3.49, which is substantially lower than the BAR levels observed in the most permissive zoning regimes in the world. For example, Singapore has many blocks allowing BAR levels in the 8 to 10 range.

Row 1 of each panel gives the total new units that are created within the city in 10 years as a result of the corresponding zoning policy; the totals exclude units created as a result of the city-wide growth trend. Row 2 shows the share of the total within city stock that the new units in row 1 represent. In Panel A, relative to a 10 year continuation of the 2004 zoning rules (Column 1), the model predicts that the 2016 zoning reform will produce approximately 47,030 net new housing units, or an approximate 1.9 percent increase in the housing stock of the city. The relatively small aggregate effect on the supply of housing is consistent with the fact the average BAR in the city only increased by 0.54, and did not change or was reduced in 48% of city blocks. In Column 3 we see that in the Double BAR reform scenario the housing stock in São Paulo increases by 27.5% relative to the 2004 zoning reform.

Row 3 gives the model-predicted average zone-level price in thousands of reais per square meter. Prices are similar in Columns 1 and 2, falling by only 0.5% on average, indicating that the 2016 reform has small impacts on average housing prices in São Paulo. Even under the Double BAR scenario we estimate only a 7.5% decrease in prices on average. One reason the price effects are small is that the suburbs are assumed to grow at an annual rate (1.7%) that is faster than the city (1%); so although the increase in housing units in the city is large relative to the stock of city housing, the supply increase relative to the total metropolitan area is smaller and therefore the price response is also commensurately smaller.[35] Another reason is that we assume each unit vacated in favor of a newly constructed unit is filled by a new migrant to the MSA, dampening downward pressure on prices. But if we assume instead that the MSA is a closed economy, as in Panel B, we see substantially larger price reductions from the 2016 and Double BAR scenarios, at 1.1% and 15.5%, respectively, even as the equilibrium unit increase is somewhat smaller.

Row 4 shows that the reform produces only a small change in the share of households living in the city versus the suburbs. This small change is unlikely to impact city neighborhood demographics. We verify this assumption in the simulation by calculating the implied demographic characteristics given the predicted demand side choice probabilities. Figure A25 compares demo-

[35] In the model, we account for differential secular trends in housing growth in the city vs. the suburbs using annualized growth rates from 2000-2010, calculated from Brazil's housing census.

graphics for city vs. suburbs under different scenarios. At baseline, the suburbs are lower income and less educated, and the 2016 zoning reform does not change the differences between city and suburbs. Only the more dramatic Double BAR reform influences neighborhood demographics by attracting more lower income and lower education residents from the suburbs to the city.

Figure 7 shows that, while there is a reasonable aggregate effect of the 2016 reform – consistent with the overall change in the housing shock, there is substantial heterogeneity across the city, mostly predicted by where BAR changed most. Figure 7 plots the distribution of zone-level changes in prices and units, and then correlates them with the average BAR change within a commuting zone. As expected, places with larger BAR shocks see more construction and lower prices. Quantitatively, the largest price reductions are roughly R$341 reais per meter squared, or around 4.6% of that zone's market price under the 2004 counterfactual. The largest supply shocks are approximately 1,549 additional units, or around 17.1% of that zone's housing stock under the 2004 counterfactual. If instead we solve for equilibrium in levels, we find that the largest price effects are around 5.4% of the zone's price under the 2004 counterfactual. These estimates are qualitatively similar to the medium-run reduced form results presented in Section 4.5.

Figure A26 maps the BAR change in the 2016 reform, as well as the model simulated predictions for changes in the number of housing units, prices, and market shares. BAR changes were larger in more outlying areas of the city, rather than in the denser core where BAR actually fell in many cases. As such, these more outlying areas saw lower prices and gained market share in a way that maps directly on to the BAR change.

5.4.2 Welfare analysis

Table 6 presents our estimates of the welfare impacts of the 2016 reform relative to a continuation of the 2004 zoning policy. These results correspond to the ″equilibrium in market shares,″ or ″open-economy″ model where we assume that demographically similar households migrate to Sao Paulo to fill new housing units. The estimates in the table give per-household average consumer surplus gains in reais at year 10 after the reform. Consumer surplus is calculated as the "inclusive value," from Small and Rosen (1981), which gives the ex-ante expected value of a utility-maximizing choice for consumer i, normalized by the individual-specific price coefficient. Consumer surplus can change after the reform for the following reasons: a) increased availability of housing options in desirable neighborhoods, allowing for a better match between households and neighborhoods; b) changes in house prices in neighborhoods that experienced housing supply shocks; c) changes in the built environment, such as the increase in the share of new units; d) changes in congestion leading to increasing commuting times.

In column (1) we update only the prices in calculating welfare in the new equilibrium, while in column (2) we update prices and the features of the built environment – the average age of buildings, average units per building, and neighborhood density – based on the supply shock. Column (3) accounts for the role of increased congestion arising from greater densification, given households' disutility of commuting time.

Focusing first on column (1), our model predicts an average increase in welfare per household of R$25.45 (approximately $7.31 USD at a 3.48 reais per dollar 2016 average exchange rate). There is an unequal distribution of surplus across demographic groups. The average college educated household has an estimated gain relative to the 2004 zoning regime of R$39.56, and the highest income quintile has a similar R$36.15 gain. High income and college households obtain the largest gains from the new housing supply for two reasons: first, they are more price sensitive in our demand model, and second, reform-induced price reductions are largest in neighborhoods where these households get the highest utility ex-ante.

Column (2) shows an average welfare gain of R$108.46, which is approximately 4.3 times larger consumer surplus gains when we allow welfare to respond to changes in the built environment. This is primarily because density increases and the age of housing units falls. These results suggest that most of the value of zoning comes through the presence of a newer housing stock and greater density, as opposed to price reductions. In Table A17 we decompose the effect of each housing characteristic and confirm that age and density are the largest contributors to these welfare gains. Interestingly, the unequal distribution of surplus across income quintiles actually narrows slightly once we incorporate changes in neighborhood characteristics, because poorer households have a greater taste for newer units. Inequality also emerges between renters and owners after updating X because owners prefer newer buildings. Column (3) shows that congestion costs erode welfare gains by 1.3% given that higher density neighborhoods increase commuting times. [36]

Table A18 presents corresponding results for our "equilibrium in levels" model, which corresponds to a "closed-economy" model where every new housing unit in the Sao Paulo MSA corresponds to a vacancy generated in the suburbs (the outside option). Under this assumption the price benefits of the reform are larger as prices fall due to less assumed demand from rural

[36] Appendix Figure A27 displays individual-level changes in expected commuting time and decomposes these changes into location and population effects.

areas; however, the aggregate welfare benefits are similar as they are mostly driven by changes in the built environment (increases in density and more new housing).

The results above referred to changes in average welfare for households in the Sao Paulo metropolitan area. The first row of Panel A of Table 7 aggregates those consumer surplus estimates to provide city-wide gains. Aggregate consumer gains correspond to R$570 million reais, or about 0.08% of the city GDP. The reform also impacted producers given the increased ability for developers to build new housing. Predicted developer profits are calculated as 10% of the total value of newly developed housing units, a conservative estimate of mark ups for real estate development. The second row of panel A shows that developer profits increase by approximately R$3.96 billion (0.57% of GDP) with the zoning reform. Those developer gains are seven times larger than the consumer welfare. These results may explain the popular view that developers benefit the most from lax development rules.

The third row of panel A of Table 7 shows a back-of-the envelope calculation of the potential effect of zoning reform on productivity. For this exercise we rely on the estimates of Glaeser and Gyourko (2018).[37] We find that future gains in productivity could be R$0.8 billion (0.11% of GDP), and are another potential justification for increasing densification. The fourth row adds up all welfare gains of the 2016 reform, which correspond to 0.76% of GDP.

In addition to the net welfare gains described above, the reform generated substantial transfers between households and landlords. Such transfers are due to nominal price changes that affected all homes in the city, not just new development. The first two rows of panel B of Table 7 show that existing homeowners and landlords face aggregated nominal housing wealth losses of R$9.88 and R$3.29 billions, respectively, due to the 2016 reform. Those losses can explain why existing real estate owners fear zoning reforms that promote more densification. However, these housing price changes do not directly impact overall city welfare, other than through the changes in consumer surplus estimated in Table 6. The intuition for this result, from Bajari, Benkard and Krainer (2005), is that to a first-order approximation, if the price of housing declines due to a zoning reform, households who are selling their homes are made worse off but other households who are buying those same less expensive houses are made better off. In equilibrium the number of buyers and sellers is identical, so welfare gains and losses offset each other. So the gains of future homeowners

[37] Glaeser and Gyourko (2018) find that a dramatic zoning reform across all cities in the United States - that allows the movement of enough workers to equalize wages across all cities - would generate a GDP gain of 2%. Assuming Sao Paulo plays an analogous role in the Brazilian economy as that of New York city in the United States economy, and also accounting for the relative magnitude of new housing developed due to the 2016 reform in Sao Paulo, we can back out potential productivity gains. See Online Appendix for details of the calculation.

and landlords mechanically corresponds to the sum of rows 1 and 2 of panel B.[38]

6 Conclusion

In this paper we study the impact of a 2016 zoning reform in São Paulo that increased the ability of developers to supply housing units by lifting limitations on permitted densification on a block-by-block basis. Using a spatial discontinuity design and the timing of the reform, we find that developers responded swiftly to obtain approximately 65 percent more permits in blocks that relaxed zoning rules. In the medium-run, the reform increased availability of homes for sale and reduced prices in neighborhoods with larger increases in allowed densification.

We also develop a framework to estimate welfare of local residents by integrating the spatial RD design in to a supply and demand model of residential housing. With this model we can estimate structural demand and supply parameters, and simulate prices, the built environment, and welfare over a ten year time horizon. Our framework accounts for both costs and benefits of densification, which allows for a more complex picture of the effects of zoning reforms.

We find mild aggregate price and supply changes due to the fact the aggregate supply response induced by the 2016 reform is small relative to the São Paulo housing stock. However, we find that certain neighborhoods that allowed for more densification had larger increases in supply and bigger reductions in prices, i.e., improved affordability. Our welfare analysis suggests that higher income and education groups benefit the most from the reform, due to the reform lowering prices in areas appealing to these groups, and their ability to move from the suburbs to areas in the city that are closer to workplaces. Welfare gains are four times larger when accounting for changes in the built environment of the city, especially with respect to neighborhood density and age of buildings. This suggests a fair amount of the economic value of zoning reforms comes through the presence of newer housing stock, as opposed to lower prices. We also find that developer profits are much larger than the realized consumer gains. Finally, such reforms negatively impact the housing wealth of existing homeowners and landlords, which may generate political backlash in the form of NIMBYism.

[38] There are caveats to the conclusion that wealth changes have no direct welfare effects. For example, if buyers and sellers face credit constraints, house price changes may differently affect the welfare of buyers versus sellers.

References

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

Bajari, Patrick, C. Lanier Benkard and John Krainer. 2005. "House Prices and Consumer Welfare." *Journal of Urban Economics* 58(3):474–487.

Balboni, Clare, Gharad Bryan, Melanie Morten and Bilal Siddiqi. 2020. "Transportation, Gentrification, and Urban Mobility: The Inequality Effects of Place-Based Policies." *Stanford University, working paper* .

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

Bayer, Patrick, Fernando Ferreira and Robert McMillan. 2007. "A Unified Framework for Measuring preferences for Schools and Neighborhoods." *Journal of Political Economy* 115(4):588–638.

Berry, Steven, James Levinsohn and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63(4):841–890.

Brueckner, Jan K and Kala Seetharam Sridhar. 2012. "Measuring Welfare Gains from Relaxation of Land-Use Restrictions: The Case of India's Building-Height Limits." *Regional Science and Urban Economics* 42(6):1061–1067.

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

Brueckner, Jan K and Somik V Lall. 2015. "Cities in Developing Countries: Fueled by Rural–Urban Migration, Lacking in Tenure Security, and Short of Affordable Housing." *Handbook of Regional and Urban Economics* 5:1399–1455.

Bryan, Gharad, Jonathan De Quidt, Tom Wilkening and Nitin Yadav. 2017. Land Trade and Development: A Market Design Approach. Technical report CEPR Discussion Paper No. DP12136.

Calder-Wang, Sophie. 2022. "The Distributional Impact of the Sharing Economy on the Housing Market." *University of Pennsylvania, working paper* .

Calonico, Sebastian, Matias D Cattaneo, Max H Farrell and Rocio Titiunik. 2017. "rdrobust: Software for regression-discontinuity designs." *The Stata Journal* 17(2):372–404.

Cattaneo, Matias D, Richard K Crump, Max H Farrell and Yingjie Feng. 2024. "On binscatter." *American Economic Review* 114(5):1488–1514.

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

Conley, Timothy G. 1999. "GMM estimation with cross sectional dependence." *Journal of econometrics* 92(1):1–45.

Davis, Morris, Jesse Gregory, Daniel Hartley and Kegon Tan. 2021. "Neighborhood Effects and Housing Vouchers." *Quantitative Economics* 12(4):1307–1346.

Diamond, Rebecca and Tim McQuade. 2019. "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development." *Journal of Political Economy* 127(3):1063–1117.

Ding, Chengri. 2013. "Building Height Restrictions, Land Development and Economic Costs." *Land Use Policy* 30(1):485–495.

Duranton, Gilles and Diego Puga. 2020. "The Economics of Urban Density." *Journal of Economic Perspectives* 34(3):3–26.

Epple, Dennis, Brett Gordon and Holger Sieg. 2010. "A New Approach to Estimating the Production Function for Housing." *American Economic Review* 100(3):905–24.

Ganong, Peter and Daniel Shoag. 2017. "Why Has Regional Income Convergence in the U.S. Declined?" *Journal of Urban Economics* 102:76–90.

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

Glaeser, Edward L, Joseph Gyourko and Raven E Saks. 2005. "Why Have Housing Prices Gone Up?" *American Economic Review* 95(2):329–333.

Goldsmith-Pinkham, Paul, Isaac Sorkin and Henry Swift. 2020. "Bartik instruments: What, when, why, and how." *American Economic Review* 110(8):2586–2624.

Gyourko, Joseph, Jonathan Hartley and Jacob Krimmel. 2021. "The Local Residential Land Use Regulatory Environment Across US Housing Markets: Evidence from a New Wharton Index." *Journal of Urban Economics* 124:103337.

Gyourko, Joseph and Raven Molloy. 2015. "Regulation and Housing Supply." *Handbook of Regional and Urban Economics* 5:1289–1337.

Harari, Mariaflavia. 2020. "Cities in Bad Shape: Urban Geometry in India." *American Economic Review* 110(8):2377–2421.

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

Kulka, Amrita, Aradhya Sood and Nicholas Chiumenti. 2022. "How to Increase Housing Affordability? Understanding Local Deterrents to Building Multifamily Housing." *Federal Reserve Bank of Boston Research Department Working Papers* 22(10).

Mullahy, John. 1997. "Instrumental-Variable Estimation of Count Data Models: Applications to Models of Cigarette Smoking Behavior." *Review of Economics and Statistics* 79(4):586–593.

Murphy, Alvin. 2018. "A Dynamic Model of Housing Supply." *American Economic Journal: Economic Policy* 10(4):243–67.

Paciorek, Andrew. 2013. "Supply Constraints and Housing Market Dynamics." *Journal of Urban Economics* 77:11–26.

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

Small, Kenneth A. and Harvey S. Rosen. 1981. "Applied Welfare Economics with Discrete Choice Models." *Econometrica* 49(1):105–130.

Song, Jaehee. 2021. "The Effects of Residential Zoning in US Housing Markets." *University of Colorado Boulder, working paper. Available at SSRN 3996483* .

Trounstine, Jessica. 2018. *Segregation by Design: Local Politics and Inequality in American Cities*. Cambridge University Press.

Tsivanidis, Nick. 2022. "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogota's TransMilenio." *University of California Berkeley, working paper* .

Turner, Matthew A, Andrew Haughwout and Wilbert Van Der Klaauw. 2014. "Land use regulation and welfare." *Econometrica* 82(4):1341–1403.

Figures

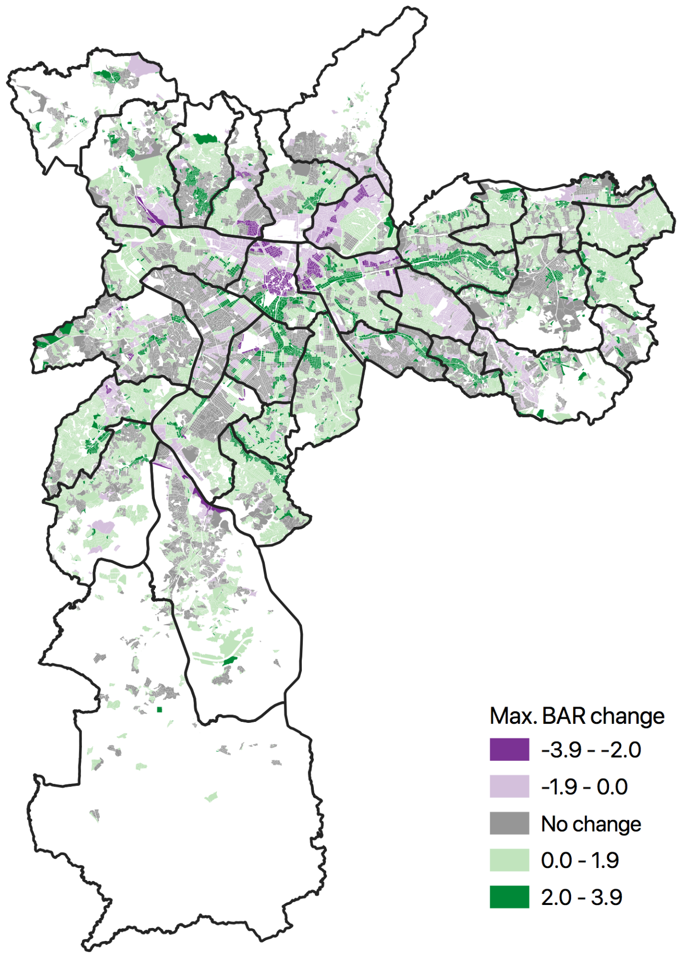

Figure 1: Map of Change in Maximum Allowed Built Area Ratio (BAR) from 2004 to 2016 Zoning Regime

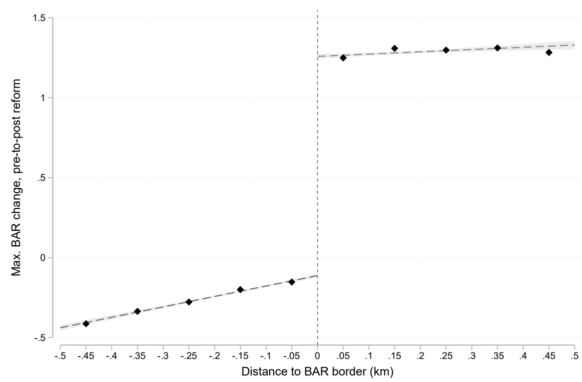

Figure 2: Built Area Ratio Change, Pre-to-Post 2016 Reform

This figure plots the change in the max BAR allowed for blocks within a .1 kilometer bin of our running variable. Control blocks are to the left of the dashed vertical line; treatment blocks are to the right. For control (treatment) blocks the running variable is the distance to the nearest treatment (control) block. A treatment block is defined as a block whose max BAR increased in the 2016 reform. Control blocks are those whose max BAR declined or stayed the same in the 2016 reform. 95% confidence intervals, in grey, are based on standard errors that are clustered at the commuting zone level. See Appendix Figure A7 for this figure using only control blocks whose max BAR stayed the same.

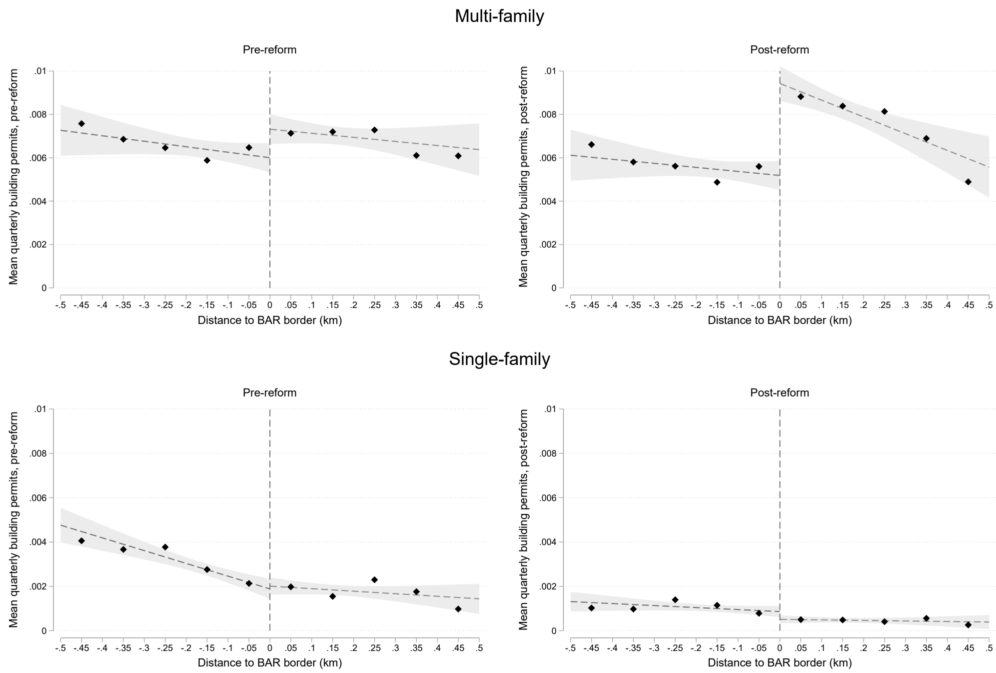

Figure 3: Multi-Family vs. Single-Family Permit Filings

This figure plots mean quarterly building permits issued for blocks within a .1 km bin of our running variable. Control blocks are to the left of the dashed vertical line; treatment blocks are to the right. For control (treatment) blocks the running variable is the distance to the nearest treatment (control) block. A treatment block is defined as a block whose max BAR increased in the 2016 reform. Control blocks are those whose max BAR declined or stayed the same in the 2016 reform. Pre-reform is 2012Q2 - 2016Q1 and post-reform is 2016Q2 - 2019Q4. 95% confidence intervals, in grey, are based on standard errors that are clustered at the commuting zone level.

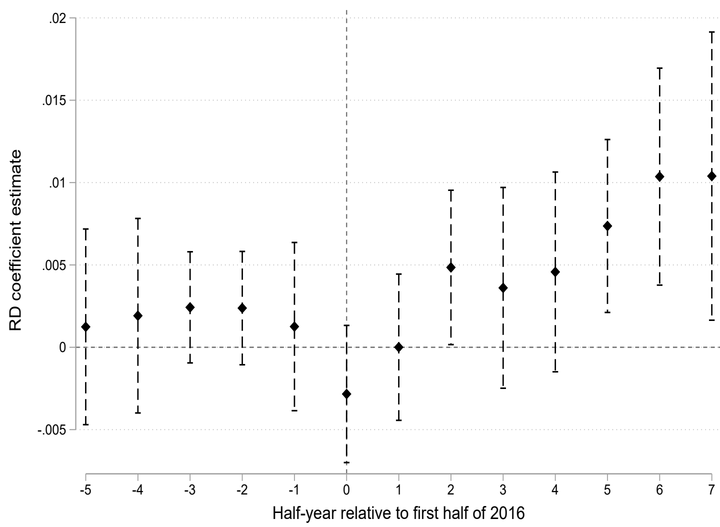

Figure 4: Dynamic RD coefficients for multifamily permit filings

This figure plots the regression-discontinuity coefficients separately estimated for half-years around the reform. The outcome variable is multifamily permit filings and the running variable is distance to the BAR border. The estimates come from a linear specification as in Column (2) of Table 2 with sub-prefeitura fixed effects. 95% confidence intervals are shown in dashed lines, calculated based on standard errors clustered at the commuting zone level.

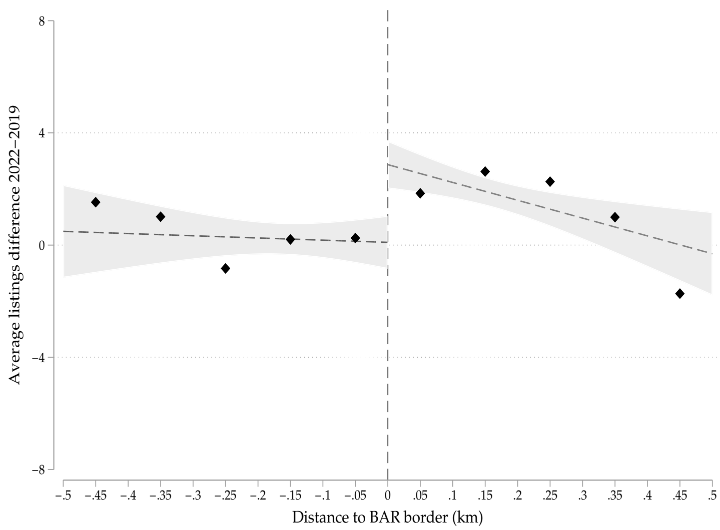

Figure 5: Listings Difference 2019-2022

This figure plots average 2019-2022 difference in number real estate of listings for sale for blocks within a .1 km bin of our running variable. Control blocks are to the left of the dashed vertical line; treatment blocks are to the right. For control (treatment) blocks the running variable is the distance to the nearest treatment (control) block. A treatment block is defined as a block whose max BAR increased in the 2016 reform. Control blocks are those whose max BAR declined or stayed the same in the 2016 reform. The light grey shading indicates 95% confidence intervals based on standard errors clustered at the commuting zone level. The mean number of listings for control blocks within 0.1 km of the BAR boundary in 2019 was 20.44; this can be compared to the treatment effect in this figure to get the proportionate effect of the zoning reform (10%). We are unable to make the same figure for the pre-treatment period (i.e. 2016-2019) due to the unavailability of granular listings data during that time period.

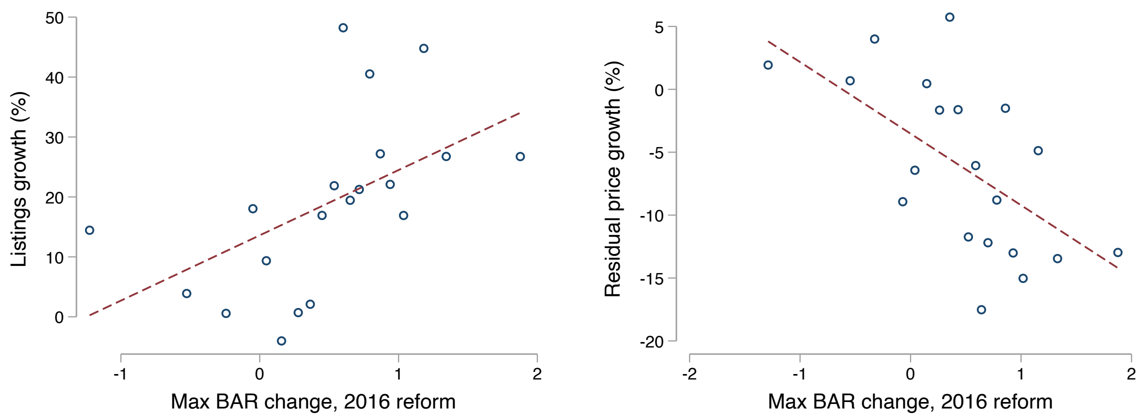

Figure 6: Growth in listings and residual prices 2022-2019

The left side figure shows a bin-scatter of commuting zone level listings growth against the average max BAR increase in the 2016 reform. The right side shows the same bin-scatter, with the y-variable changed to growth in average residual price within the commuting zone. Growth rates are for 2022 relative to 2019. Dashed lines show estimated slopes based on difference-in-difference models of the respective outcomes (changes in listings and residual prices) on max BAR changes presented in Table A9.

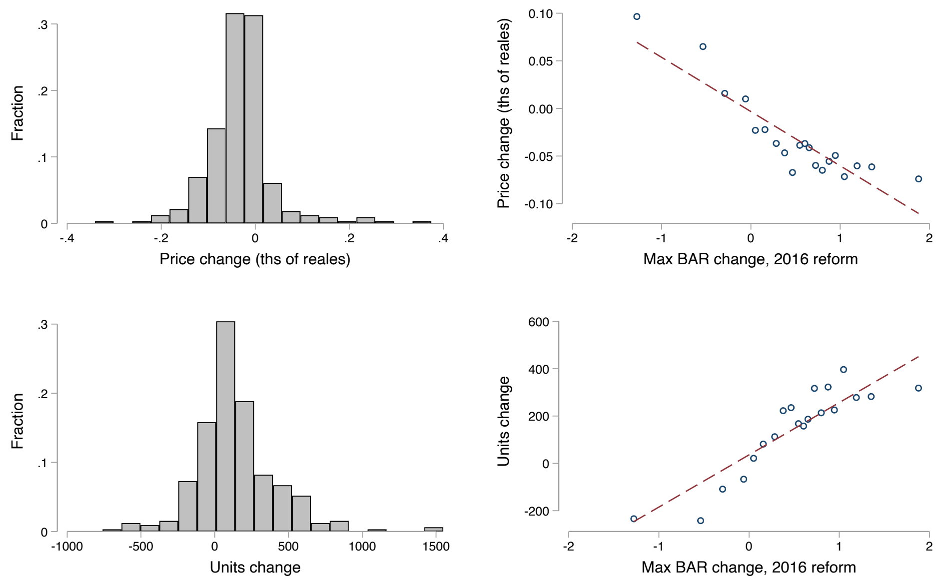

All changes relative to 2004 reform

Figure 7: Equilibrium changes under the 2016 reform

The top-left figure shows the distribution, across commuting zones, of the 2016 reform model simulated price changes for 2026 relative to a continuation of the 2004 zoning regime until 2026. The bottom-left figure shows the same for number of housing unit changes. The top-right figure shows a bin-scatter of the 2016 reform model predicted commuting zone level price change against the average max BAR increase in the 2016 reform. The bottom-right shows the same bin-scatter, with the y-variable changes to increase in number of housing units within the commuting zone.

Tables

Table 1: RD first stage

Outcome	Max BAR change			
	(1)	(2)	(3)	(4)
Panel A: No sub-prefeitura FE				
Treat BAR	1.519***	1.426***	1.365***	1.354***
	(0.029)	(0.029)	(0.031)	(0.031)
Specification	Base	Linear	Quadratic	Cubic
Observations	43231	43231	43231	43231
Mean of Dep. Variable	-0.153	-0.153	-0.153	-0.153
Panel B: With sub-prefeitura FE				
Treat BAR	1.516***	1.453***	1.386***	1.355***
	(0.025)	(0.027)	(0.029)	(0.030)
Specification	Base	Linear	Quadratic	Cubic
Observations	43225	43225	43225	43225
Mean of Dep. Variable	-0.153	-0.153	-0.153	-0.153

Standard errors clustered by commuting zones in parentheses. Specification refers to the order of the polynomial for the running variable, which is distance to the RD boundary. The polynomial is always interacted with the treatment indicator. Sample is all city blocks with zoning information. Mean of dependent variable is the change in max BAR pre- versus post- calculated for control blocks within 0.1 km of the BAR boundary, indicating a .15 max BAR decline for blocks just to the left of the cut-off. *$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 2: RD reduced form

Outcome	New multi-family building permits			
	(1)	(2)	(3)	(4)
Panel A: No sub-prefeitura FE				
Treat BAR	0.00170*	0.00409***	0.00481***	0.00428***
	(0.00071)	(0.00073)	(0.00078)	(0.00086)
Specification	Base	Linear	Quadratic	Cubic
Observations	43231	43231	43231	43231
Mean of Dep. Variable	0.00560	0.00560	0.00560	0.00560
Panel B: With sub-prefeitura FE				
Treat BAR	0.00198***	0.00228***	0.00268***	0.00313***
	(0.00052)	(0.00056)	(0.00066)	(0.00080)
Specification	Base	Linear	Quadratic	Cubic
Observations	43225	43225	43225	43225
Mean of Dep. Variable	0.00560	0.00560	0.00560	0.00560

Standard errors clustered by commuting zones in parentheses. Specification refers to the order of the polynomial for the running variable, which is distance to the RD boundary. The polynomial is always interacted with the treatment indicator. Sample is all city blocks with zoning information. Mean of dependent variable calculated for control blocks within 0.1 km of the BAR boundary. *$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 3: Supply estimates: Poisson IV regressions

Outcome	All new buildings			Single	Multi
	(1)	(2)	(3)	(4)	(5)
Max BAR	0.386***	0.789***	0.747***	-0.048	0.870***
	(0.099)	(0.166)	(0.163)	(0.335)	(0.188)
Price	0.140***	0.149***	0.415***	0.425*	0.387**
	(0.034)	(0.035)	(0.101)	(0.214)	(0.139)
Density	0.222	0.116	-0.292	-0.176	-0.251
	(0.118)	(0.125)	(0.195)	(0.435)	(0.247)
Age	0.017*	0.016*	-0.019	-0.045	-0.023
	(0.008)	(0.008)	(0.015)	(0.035)	(0.020)
Units per building	-0.005	-0.006	-0.013	-0.025	-0.023*
	(0.004)	(0.005)	(0.007)	(0.014)	(0.010)
Historical preservation	-0.747*	-0.699	-0.453	-0.009	-0.831
	(0.347)	(0.358)	(0.387)	(0.788)	(0.458)
Q	1.764e-29	1.658e-29	2.684e-29	2.712e-29	1.007e-28
Observations	1182	1182	1182	1182	1182
IVs	None	RD	RD, Bartik	RD, Bartik	RD, Bartik

Robust standard errors in parentheses. Results are from the estimation of a fuzzy regression discontinuity (RD) exponential (Poisson) model, estimated with GMM, on the sample of subprefeitura-quantiles. The RD treatment indicator instruments for Max BAR, while the Bartik labor demand shock instruments for price. All models use an additive error specification to form moment conditions. All specifications include controls for the running variable interacted with the treatment, and the following zoning parameters: maximum shadow ratio, minimum and basic BAR of 2004 and 2016, max BAR of 2004, maximum height, min and max. front setback and maximum area of 2016, (zoning variables averaged within subprefeitura-quantile). Q-statistic gives the value of the GMM criterion function at the optimal parameters. The outcome variable is the number of total new building, single-family, or multi-family permit applications between 2016-2019, as indicated. *$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 4: Estimated demand parameters

Demographic	Price (1)	Travel time (2)	RCMA (3)	Age (4)	Units (5)	Density (6)	Paved (7)	Income (8)	Education (9)
Household size	-0.087	0.127	0.026	-0.064	-0.071	-0.050	0.002	-0.042	-0.041
	(0.027)	(0.012)	(0.023)	(0.018)	(0.040)	(0.022)	(0.013)	(0.031)	(0.038)
Age	-0.064	-0.153	-0.044	0.056	-0.151	0.005	0.035	0.060	0.123
	(0.026)	(0.011)	(0.024)	(0.019)	(0.029)	(0.021)	(0.017)	(0.033)	(0.039)
Renter	-0.060	-0.122	0.004	0.198	-0.002	0.034	0.037	-0.003	-0.015
	(0.026)	(0.011)	(0.022)	(0.019)	(0.027)	(0.018)	(0.017)	(0.032)	(0.036)
Income	-0.077	-0.053	0.039	0.110	-0.145	0.081	0.115	0.467	-0.005
	(0.034)	(0.023)	(0.033)	(0.036)	(0.060)	(0.024)	(0.046)	(0.041)	(0.058)
College degree	-0.185	-0.047	-0.017	-0.011	-0.163	0.068	0.058	-0.314	0.880
	(0.036)	(0.014)	(0.031)	(0.026)	(0.050)	(0.025)	(0.030)	(0.037)	(0.046)
Base coefficients	-1.975	-2.413	0.749	-0.980	-0.647	0.765	0.090	-0.420	0.572
	(0.454)	(0.021)	(0.225)	(0.124)	(0.432)	(0.159)	(0.104)	(0.250)	(0.294)

Results are from the estimation of demand-side preference parameters using two-step maximum likelihood and 2SLS. Top row gives variable names, while left-most column gives the demographic variables. Estimation sample is 329 commuting zones and 24,800 individual households. All location characteristics including price are standardized relative to the zone-level sample mean and standard deviation. Travel time is normalized across all individual-zone combinations. Base coefficients are from column (9) of Table A16, which instruments for housing prices using the average spatial and housing characteristics of zones 5-20 miles from a zone centroid. These characteristics are favela share of zone area, slope, RCMA, and housing stock age. Bootstrapped standard errors with 500 replications in parentheses.

Table 5: Simulation results: Housing prices and quantities

Scenario	2004 zoning	2016 zoning	Double BAR
Max BAR	1.55	2.09	3.49
Panel A: Equating shares			
New units (ths)	166.357	213.387	844.607
New units (share of stock)	0.068	0.087	0.343
Avg. price (ths of reales)	6.114	6.083	5.655
Inside share	0.587	0.591	0.637
Panel B: Equating levels			
New units (ths)	157.931	199.714	667.027
New units (share of stock)	0.064	0.081	0.271
Avg. price (ths of reales)	5.987	5.922	5.060
Inside share	0.606	0.615	0.713

Table shows zone-level results from equilibrium simulations over a ten year period from 2016 to 2026, under three different zoning scenarios, as indicated in table header. Double BAR scenario holds BAR constant at 2004 levels for all locations where BAR was reduced in 2016, and doubles the post-2016 BAR value in all locations where BAR was increased in 2016. First row shows average block-level maximum allowable BAR under each scenario. Panel A equates market shares in the equilibrium condition, implicitly assuming that all new construction within the city is occupied by new migrants from outside the MSA. Panel B equates levels in the equilibrium condition, implicitly assuming that all new construction in the city is occupied by migrants from outside the city but within the MSA. Average price is thousands of reales per square meter. Inside share is fraction of households living within the municipality.

Table 6: 2016 zoning reform simulation results: Household consumer surplus

Update	P	X	τ
	(1)	(2)	(3)
Average consumer gain	25.45	108.46	107.02
By demographic group			
Owner	25.08	115.94	114.54
Renter	26.55	86.04	84.47
Non-college	21.59	94.18	92.98
College	39.56	160.68	158.36
By income quintile			
1	19.69	86.06	84.99
2	21.31	94.76	93.52
3	23.38	102.18	100.85
4	27.33	117.97	116.40
5	36.15	143.77	141.72

Table shows expected change in individual consumer surplus over a ten year period from 2016 to 2026, measured in Brazilian reais per household (approximately 5 reais to the USD during this time period), from equilibrium simulation of the 2016 zoning reform for different subgroups. Column (1) updates only equilibrium prices from the 2016 reform scenario, while column (2) updates both prices and the housing and neighborhood attributes included in X_j. Column (3) updates all variables, including travel time τ. All changes are evaluated relative to 2004 (status quo) zoning.

Table 7: 2016 zoning reform aggregate welfare estimates and wealth transfers

	R$ bi	% of GDP
Changes in welfare	(1)	(2)
Consumer gains	0.57	0.08
Developer profits	3.96	0.57
Productivity gains	0.80	0.11
Total	5.33	0.76
Wealth transfers from current to future real estate owners		
Homeowners	9.88	1.41
Landlords	3.29	0.47
Total	13.17	1.88

The first panel shows calculations for aggregate policy effects on welfare changes based on the 2016 reform equilibrium simulation over a ten year period from 2016 to 2026. All changes in welfare are evaluated relative to the 2004 status quo zoning. GDP refers to 2017 total output of Sao Paulo city. The second panel reports aggregate changes in nominal house prices due to the 2016 reform. Homeowners refers to value of all owned-occupied units, while landlords refers to value of all rental units. See text for more details.