Zoning: Externalities or Misallocation?*

Yu-Hsin Ho
National Taiwan University

Chang-Tai Hsieh
University of Chicago

Wen-Tai Hsu
Academia Sinica

Yu-Jhih Luo
Washington University

March 26, 2026

This paper provides an empirical analysis of the effects of residential versus commercial zoning in U.S. cities. We exploit an institutional contrast with Taipei, where strict separation between residential and commercial uses is largely absent. Using parcel-level data, we compare the allocation of space between housing and production across neighborhoods in Taipei and U.S. metropolitan areas. Guided by a model in which neighborhoods differ in productivity and residential amenities, we interpret differences across neighborhoods in land-use patterns as reflecting zoning-induced wedges between residential and commercial space. We find substantially greater dispersion in the ratio of residential space per resident to commercial space per worker in U.S. cities relative to Taipei. Counterfactual estimates suggest that reducing this dispersion to Taipei's level would raise welfare by 15.2 percent on average. Zoning also leads to more specialized neighborhoods, with neighborhoods in U.S. cities far more likely to be predominantly residential or commercial. This pattern could reflect the internalization of negative spillovers from the co-location of workers and residents. However, if mixing workers and residents generates positive spillovers, zoning instead exacerbates inefficiencies by producing excessively segregated neighborhoods.

* We thank Enrico Moretti for enormously helpful discussions.

1 Introduction

This paper studies the economic consequences of separating residential vs. commercial land uses through zoning. While separation between housing and employment activities is a defining feature of land-use regulation in most U.S. cities, many cities abroad operate without this type of land-use regulation. For instance, cities in Japan and Taiwan do not enforce strict separations between residential and commercial zones. We exploit this institutional contrast to identify the effects of zoning in the U.S.

Using detailed parcel-level data, we examine how space is allocated between residential and commercial uses across neighborhoods in Taipei and in American cities. In Taipei, where land-use zoning is absent, land-use patterns are shaped entirely by market forces. In contrast, allocations in American cities reflect both market dynamics and zoning regulations. We interpret the differences in observed land-use patterns between Taipei and U.S. cities as evidence of the effects of zoning.

To guide our analysis, we use a model of a city with heterogeneous neighborhoods that differ in terms of their productivity and residential amenities, with space serving both production and housing purposes. In the absence of zoning, the distribution of residential and commercial space across neighborhoods is driven entirely by comparative advantage in production vs. amenities across neighborhoods. We model zoning as a wedge between the price of commercial and residential space, which distorts this allocation.

We use our parcel level data to infer this price wedge by examining differences between residential space per resident to commercial space per worker across neighborhoods. Our approach relies on the assumption that the marginal product of space is proportional to its average product. To account for the possible specification error, we benchmark these gaps against those observed in Taipei, where zoning is absent. Specifically, we interpret the *additional* dispersion in U.S. cities, relative to Taipei, as reflecting the price wedge induced by American zoning.

We report three main findings. First, the dispersion across neighborhoods in residential space per resident relative to commercial space per worker is significantly larger in U.S. cities compared to Taipei. The 90-10 gap in this ratio is 1.5 log points higher in the average American city compared to Taipei. Even in the least dispersed American city in our sample, Austin, the 90-10 gap remains nearly 1 log point higher than in Taipei.

Second, we estimate the welfare costs of zoning by considering a counterfactual in which the dispersion in land-use distortions is reduced to the level observed in Taipei. Under this scenario, average aggregate welfare in U.S. cities increases by 15.2 percent. The gains vary widely, from 5.4 percent in Austin to as much as 39.4 percent in Chicago.

Third, zoning leads to greater specialization of neighborhoods in U.S. cities. Empirically, American neighborhoods are far more likely to be predominantly residential or commercial compared to those in Taipei. Consistent with this, our model predicts that, absent zoning, U.S. cities would exhibit more mixed-use neighborhoods with a more balanced distribution of activities.

This last result is central to interpreting the role of zoning in addressing externalities. Although zoning may impose economic costs, it could also generate benefits by internalizing

externalities. To explore this, we extend our model in two directions.

First, we introduce negative externalities from the co-location of commercial activity and residents, which create incentives to separate land uses. In this framework, zoning can reduce such externalities by promoting spatial segregation. While U.S. cities do exhibit greater segregation, the pattern of zoning appears misaligned with efficiency: residential zoning is more prevalent in neighborhoods with a comparative advantage in production, rather than in those with stronger residential amenities. Thus, although zoning may mitigate negative externalities (to the extent that they are present), it does so in the most costly way possible.

Second, we analyze a model with positive externalities from mixing residents and workers, motivated by the idea that local services depend on activity throughout the day. In such a setting, optimal policy would promote mixed-use neighborhoods by introducing residential zoning in areas with a comparative advantage in production, and conversely encouraging employment activity in areas that otherwise would be primarily residential. This pattern is consistent with what we observe in the data and tends to increase mixed-use development.

However, we also find that a large part of the zoning does not correspond to the comparative advantages of the neighborhoods. These mismatches generate more spatial segregation between residential and employment areas. Empirically, this second effect – zoning that is poorly aligned with local comparative advantage – is much larger than the first. As a result, the overall effect of zoning in U.S. cities is to reduce, rather than promote, mixed-use neighborhoods.

The remainder of the paper proceeds as follows. We begin by presenting the model and its implications for the spatial allocation of land and activity. We then describe the parcel-level data for Taipei and U.S. cities, followed by a discussion of zoning institutions in Taipei and evidence supporting its treatment as effectively unzoned. Next, we document patterns of neighborhood specialization and estimate the dispersion in the marginal product of land across uses. We then use these estimates to quantify the welfare costs of zoning. The final section extends the analysis to incorporate externalities and evaluates the effectiveness of zoning in addressing them.

2 Effect of Zoning on Space and Labor Allocation

This section introduces residential vs. commercial zoning in a model of a city with neighborhoods that differ in productivity and residential amenities. The goal is to show how zoning affects the allocation of land and people across neighborhoods.

2.1 Environment

The production technology in neighborhood i is $Y_i = A_i H_{Ci}^{\alpha} L_{Ei}^{1-\alpha}$ where H_{Ci} and L_{Ei} denote space and effective labor. Cost minimization implies that the relative demand

for labor and commercial space is:

$$\frac{H_{Ci}}{L_{Ei}} = \frac{1-\alpha}{\alpha}\frac{w_i}{q_{Ci}} \tag{1}$$

where w_i and q_{Ci} denote the local wage and the price of commercial space. Using this condition, and normalizing the price of output to 1, the dual of the production technology is:

$$w_i^{\alpha} = \frac{A_i}{q_{Ci}^{1-\alpha}}$$

where w_i is the wage and q_{Ci} is the price of commercial space.

Labor income of a worker who lives in neighborhood i and works in neighborhood j is $w_j t_{ij} \mu$ where $t_{ij} \leq 1$ is an iceberg commuting cost and μ_j is the worker's idiosyncratic productivity from working in location j. We assume μ_j is an iid Frechet distribution with shape parameter ϵ. The expected income of a person living in i is

$$x_i \equiv \Gamma\left(\frac{\epsilon-1}{\epsilon}\right) W_i \, + \, r.$$

where $W_i \equiv \left(\sum_j (w_j \, t_{ij})^{\epsilon}\right)^{1/\epsilon}$.[1] The first term denotes expected wage income; the second term denotes property income.[2]

Utility is $B_i C_i^{1-\gamma} H_{Ri}^{\gamma}$ where C_i is consumption of the output good and H_{Ri} is residential land. Utility maximization implies the following demand for residential land per resident:

$$\frac{H_{Ri}}{L_{Ri}} = \gamma \frac{x_i}{q_{Ri}} \tag{2}$$

where q_{Ri} is the price of residential space. Expected indirect utility of a worker in i is then

$$V_i = \frac{B_i \, x_i}{q_{Ri}^{\gamma}}$$

We model zoning as a price wedge between the price of residential vs. commercial space:

$$\tau_i \equiv \frac{q_{Ci}}{q_{Ri}}$$

$\tau_i > 1$ denotes pro-residential zoning and $\tau_i < 1$ is pro-commercial zoning. Note that τ_i also potentially includes the effect of building restrictions as well as the effect of commercial vs. residential zoning. However, for building restrictions to appear as τ_i, two conditions have to hold. First, building restrictions must have a differential effect

[1] $\Gamma()$ is the gamma function.

[2] This formulation assumes property income is the same for all people regardless of where they live. Think about a mutual fund that owns all properties in the city and all residents of the city own equal shares of the mutual fund. This assumption guarantees that market allocations are welfare maximizing in the absence of externalities (see Gaubert and Fajgelbaum).

on commercial vs. residential spaces. Second, this differential effect also must differ between neighborhoods in a city.

Finally, we assume people freely move across neighborhoods until utility is the same in all neighborhoods. This assumption implies that the local wage can be written as:

$$w_i^{\alpha} \propto \frac{A_i}{\tau_i^{1-\alpha} B_i^{(1-\alpha)/\gamma}} \tag{3}$$

The local wage is increasing in local productivity A_i and decreasing in local amenities B_i and the extent of pro-residential zoning τ_i.

2.2 Equilibrium

In this subsection we solve for the distribution of employment, residents, and space across neighborhoods.

We start with the distribution of residential vs. commercial space. Combining the first order conditions for production and utility, we get:

$$\frac{H_{Ri}}{H_{Ci}} = \frac{\gamma\,\alpha}{1-\alpha}\,\frac{x_i}{w_i}\,\tau_i\,\frac{L_{Ei}}{L_{Ri}} \tag{4}$$

The left hand side is the ratio of residential to commercial space. The terms on the right hand side are the Cobb-Douglas parameters in production and utility (first term), the ratio of income to wages in the neighborhood (second term), the price gap between residential and commercial space (third term), and the employment share vs. the residential share of the neighborhood (fourth term).

Next we solve for the distribution of employment. Assuming that each worker chooses a work location that maximizes their indirect utility, the probability that a resident of k chooses to work in i is $\left(\frac{w_i\,t_{ki}}{W_k}\right)^{\epsilon}$. The employment share of location i (in effective units) is then:

$$L_{Ei} = w_i^{\epsilon-1} \sum_k \frac{t_{ki}^{\epsilon}}{W_k^{\epsilon-1}} L_{Rk} \tag{5}$$

In the absence of commuting costs, the employment share is proportional to $w_j^{\epsilon-1}$. Referring back to the equation for the local wage (3), the employment share of the neighborhood is increasing in its comparative advantage in production relative to residential amenities and in the extent of pro-commercial zoning (lower τ), with an elasticity of $\frac{\epsilon-1}{\alpha}$.

The residential share of a neighborhood is proportional to:[3]

$$L_{Ri} = B_i^{1/\gamma}\, x_i^{(1-\gamma)/\gamma}\, H_i\, (1 + H_{Ci}/H_{R_i})^{-1} \tag{6}$$

In the absence of commuting cost, expected income x_i is the same in all neighborhoods. In this case, the residential share of a neighborhood is increasing in local amenities B_i

[3]The complete equation for the residential share is the ratio of the right hand side of equation (6) relative to its sum across all neighborhoods $\sum_k B_k^{1/\gamma}\, x_k^{(1-\gamma)/\gamma}\, H_k\, (1 + H_{Ck}/H_{R_k})^{-1}$.

and total space H_i. Furthermore, since H_{C_i}/H_{Ri} is increasing in A_i and decreasing in B_i and τ_i, the residential share is also increasing in τ_i and decreasing in local productivity A_i.[4]

In sum, in the absence of zoning, the distribution of employment and residents is entirely driven by the distribution of A/B across neighborhoods. Neighborhoods with a comparative advantage in production (high A/B) are more specialized in production. There are more workers than residents, and more space is allocated for commercial uses compared to residential uses. On the other hand, neighborhoods with a comparative advantage in residential amenities (low A/B) are more specialized in residents. More space is allocated for residential use, and there are more residents than workers.

With zoning, the allocation of space between commercial and residential use and the ratio of workers to residents is not only driven by the distribution of A/B. Pro-residential zoning ($\tau > 1$) increases the share of residents and lowers the share of workers; pro-commercial zoning ($\tau < 1$) results in more workers and less residents.

Finally, in the full model with commuting costs, the residential share also depends on local income x_i which is increasing in the neighborhood's proximity to neighborhoods with high wages. In addition, local employment also depends on its proximity to populated neighborhoods.

2.3 Inference

The forcing variables in the model are $\{A_i\}$, $\{B_i\}$, and $\{\tau_i\}$. In this subsection, we describe how we will infer these variables from data on the distribution of employment, residential population, and commercial vs. residential space across neighborhoods.

To make clear the intuition behind the inference exercise, consider the model without commuting cost. We proceed in four steps.

First, we infer employment in effective units L_{Ei} from data on the share of workers L_{Ci}:

$$L_{Ei} = L_{Ci}^{\frac{\epsilon-1}{\epsilon}}.$$

Note that $\frac{\epsilon-1}{\epsilon} < 1$, which implies that the elasticity of effective employment with respect to the number of workers is less than 1. This reflects the effect of selection in the model: the productivity of an average worker declines as more workers choose the neighborhood.

Second, we infer τ from the ratio of space to people for residential vs. commercial use:

$$\frac{H_{Ri}/L_{Ri}}{H_{Ci}/L_{Ei}} = \frac{\gamma\,\alpha}{1-\alpha}\,\frac{1}{w_i}\,\tau_i$$

where we infer w_i from $L_{Ei} \propto w_i^{\epsilon-1}$. Intuitively, pro-residential zoning ($\tau > 1$) increases the land to labor ratio for residential use compared to commercial space, and vice-versa. The key assumption is that the Cobb-Douglas parameters $\frac{\gamma\,\alpha}{1-\alpha}$ are the same in all neighborhoods.[5]

[4] Equations (4) and (5) imply $\frac{H_{Ci}}{H_{Ri}} \propto \frac{L_{Ri}}{\tau_i\, w_i^{\epsilon-2}}$.

[5] We pick $\frac{\gamma\,\alpha}{1-\alpha}$ in each city such that the weighted average of $\tau_i = 1$.

Third, we infer the distribution of local amenities B_i from data on residential population and space use. Specifically, the equation for residential population (6) can be expressed as:

$$B_i \propto \left(L_{Ri}\left(1 + H_{Ci}/H_{Ri}\right) H_i^{-1}\right)^{\gamma}$$

The intuition is that, conditional on the allocation of space in the neighborhood, local resident population L_{Ri} is increasing in the value of local amenities B_i: better neighborhoods attract more residents.

Fourth, we rewrite the equation for the employment distribution (5) as:

$$A_i \propto L_{Ei}^{\frac{\alpha}{\epsilon-1}} \tau_i^{1-\alpha} B_i^{(1-\alpha)/\gamma}$$

The intuition is that, conditional on τ and B_i, neighborhoods with a larger share of employment are more productive (higher A_i).

In the full model with commuting costs, we need to take two additional things into account. First, employment in a neighborhood also depends on its distance to populated residential neighborhoods. Second, local income x_i also varies across neighborhoods depending on each neighborhood's distance from high wage neighborhoods, and this variable also matters for the distribution of residential population and for the inference of τ from the space/labor ratio for commercial vs. residential use.[6]

3 Data

This section describes the data used to calculate H_{Ri}, H_{Ci}, L_{Ei}, and L_{Ri} in each neighborhood in U.S. cities and in Taipei.

We start with U.S. cities. We focus on 36 MSAs with populations greater than a million. We define a neighborhood in a MSA as census block group. We restrict our attention to census block groups that have a population density that exceeds 1,000 residents or an employment density that exceeds 500 workers per square kilometer.[7] The average MSA has 1600 block groups, with an average size of 0.76 kilometers (or about 6 square blocks).

Our key data are parcel-level data on all properties in the U.S. provided by Cotality, a major real-estate data provider. The backbone of this data are local tax records complemented with property appraisals (from lenders and insurance companies). We use Cotality's information on land use, floor area of the property, and the geographic coordinates of the property. We remove land used for public administration, education, recreation, transportation, warehouses, and industrial uses. We then aggregate the parcel level data into total floor area for commercial and residential use in each census block group.

Lastly, we obtain the residential population in each census block group from the American Community Survey. We obtain the employment population in each census

[6]The appendix details the inference exercise in the full model with commuting costs.
[7]This restriction lowers the average number of neighborhoods by 16%.

block group from the LEHD's Origin-Destination Employment Statistics (LODES). We drop employment in the same industries that we removed in the floor space data.

Turning to Taipei, the smallest administrative units are indeed called "neighborhoods." There are about 1600 neighborhoods in Taipei, and they are roughly the same size as a U.S. census block group. In each neighborhood, we calculate the total floor space used for residential and commercial use. We obtain this data from the Ministry of Finance's tax records in 2015. The tax data provides each property's location, floor area, and the use of the property. Since it is common for a property to have multiple uses, the tax records provide information on the floor space associated with each use. Since each use is charged a different tax rate, it is important for the Ministry of Finance to get this information correct. After imposing the same restrictions we used on the U.S. parcel level data, we aggregate the floor space in Taipei into commercial and residential use in each of Taipei's neighborhoods.

We obtain the residential population in each Taipei neighborhood from the Household Registration Records of the Ministry of the Interior. Our employment data is from the micro-data of the 2016 Economic Census, which is a complete enumeration of all establishments conducted every five years. The Economic Census provides the location and the four digit industry of each establishment, from which we compile total employment in each neighborhood.[8]

4 Flexible Zoning in Taipei

In this section, we argue that Taipei essentially does not have binding zoning separating residential and commercial uses. Establishing this institutional fact is important because we will use Taipei as a benchmark for our estimates of τ.

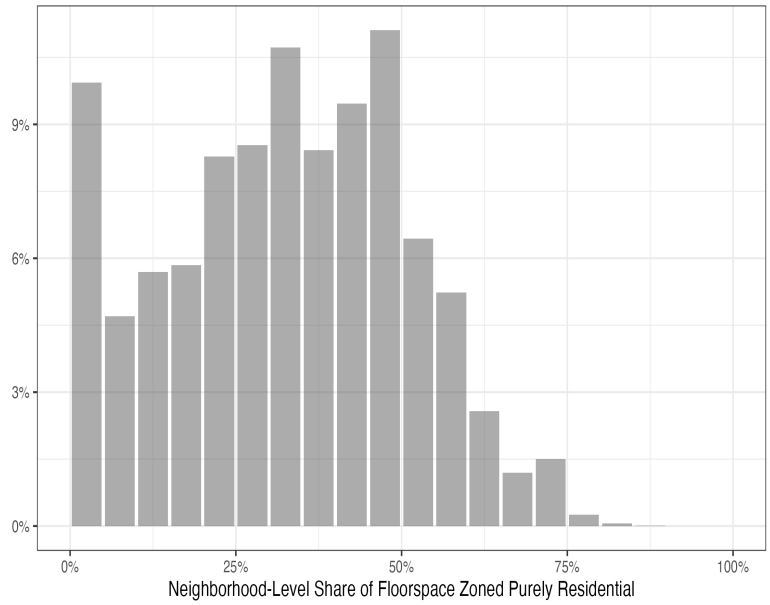
To begin, Taipei has no commercial zoning. It does have residential zoning, but this type of zoning is at the parcel level, not at the neighborhood level. We obtained property level data from the Departments of Land Administration in Taipei and New Taipei on the zoning of each parcel.[9] It is also common for the zoning to differ by floor of a given building. For example, the R3 residential code – the most common zoning code in Taipei – zones all floors above the second floor as residential use only. The first and second floors are zoned as "flexible use" which means it can be used for any purpose. Figure 1 shows the distribution of the share of residential zoned land in Taipei neighborhoods compiled from the zoning data. The figure also shows that residential zoned space is dispersed across neighborhoods in Taipei. There are essentially no neighborhoods where more than 75% of the space is zoned residential. In total, about 30% of total floor space in Taipei are zoned for residential use only. The remainder, or 70%, is zoned as "flexible use."

Figure 2 shows that residential zoning in Taipei is almost never binding. Each observation is a Taipei neighborhood. The figure plots the actual space used for residential purposes (on the y-axis) in the neighborhood against the space zoned as residential only

[8]We drop establishments in the same industries we dropped in the U.S. data.

[9]Taipei is administratively divided into two cities: Taipei and New Taipei.

Figure 1: Distribution of Residential Zoned Space in Taipei Neighborhoods



Notes: Figure shows the distribution of the share of space zoned for residential use only in each neighborhood in Taipei

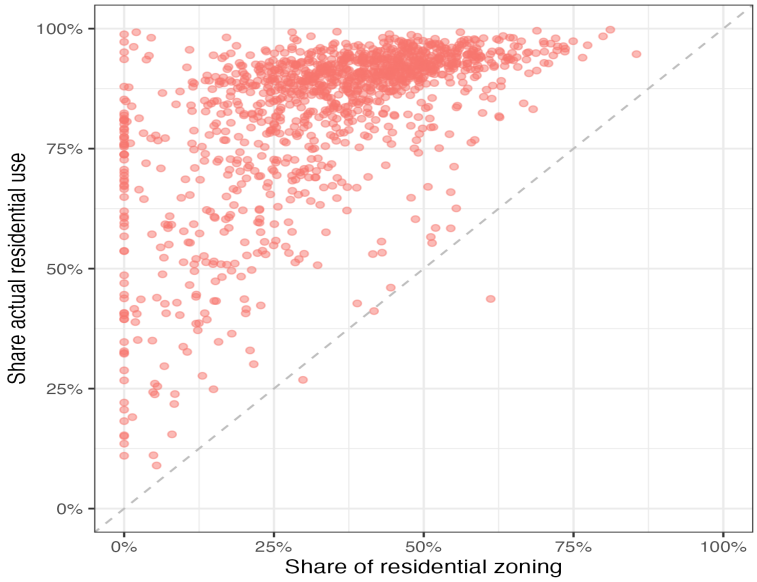
(on the x-axis). If residential zoning were binding, the observations would lie exactly on the 45 degree line. It is obvious that all the observations are significantly above the 45 degree line. It is also clear that there is a wide dispersion in land use among neighborhoods with a similar amount of residential zoned space. For example, among the neighborhoods where none of the land is zoned residential, the actual share of land used for residential purposes ranges from 10% to almost 100%.

Figure 3 illustrates the flexibility in the use of space in Taipei vs. U.S. cities by showing the changes over time in the share of residential space. Specifically, the figure shows the distribution of the change in the share of residential space in a neighborhood over five years in Taipei vs. the average American city. The dispersion of the change in the share of residential change of a neighborhood is much larger in Taipei vs. U.S. cities. In the typical American city, about 40% of the neighborhoods see zero change in the share of space used for residential purposes. In contrast, in Taipei the residential share increases in many neighborhoods and falls in other neighborhoods.

5 Specialization in American Cities

In this section, we use data on the allocation of space between residential and commercial use and the allocation of workers and residents across neighborhoods to infer two driving

Figure 2: Residential Use vs. Residential Zoning in Taipei



Notes: Figure shows the percentage point change in the share of residential space in a neighborhood over five years. U.S. is average of 36 large MSAs.

forces in the model: (1) the distribution of the price wedge between residential and commercial uses τ across neighborhoods; (2) the distribution of comparative advantage A/B, also across neighborhoods.

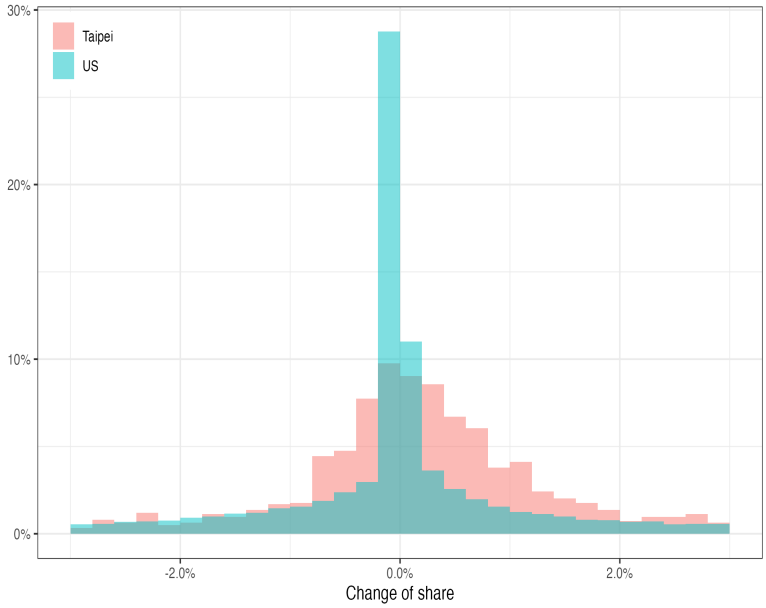
In the model, the specialization of neighborhoods, both in terms of commercial vs. residential space and in terms of workers vs. residents, is driven by comparative advantage in production vs. amenities (A/B) and by the price wedge between residential and commercial space τ. We therefore begin by showing the net effect of these two forces on the extent of specialization of space and workers vs. residents.

We begin by visualizing the specialization of neighborhoods in commercial vs. residential space in Figure 4 (left panel). The figure plots the distribution of $H_{Ri}/H_R - H_{Ci}/H_C$ of each neighborhood, normalized so that the maximum of this index is 1 and the minimum is -1.[10] This index is at a maximum when all the space in the neighborhood is residential space. Its minimum is when all the space in the neighborhood is used for commercial purposes. The index is zero when space is equally shared between commercial and residential use.

The figure clearly shows the segregation of land use (commercial vs. residential) between neighborhoods in American cities. Note the mass of the distribution at $+1$.

[10] When the index is positive, it is normalized by the ratio of the total residential space in the city to the total space in the neighborhood. When the index is negative, it is normalized by the ratio of the total commercial space in the city to the total space of the neighborhood.

Figure 3: Distribution of Change in Residential Share of a Neighborhood



Notes: Figure shows the percentage point change in the share of residential space in a neighborhood over five years. U.S. is average of 36 large MSAs.

These are neighborhoods where the space is 100% residential. The contrast with Taipei is remarkable, not only at the right tail but also on the left tail of the distribution.

The right panel in Figure 4 shows the distribution of the specialization of neighborhoods in terms of workers vs. residents. The figure plots the distribution of $L_{Ri}/L_R - L_{Ci}/L_C$ of each neighborhood, again normalized so that the maximum of this index is 1 and the minimum is -1.[11] When a neighborhood only has residents the index is 1, and when a neighborhood only has workers the index is -1. And when a neighborhood has the same number of workers and residents, the index is zero.

Here again the specialization of neighborhoods in American cities, this time in terms of workers vs. residents, is quite remarkable. The bimodality of the index is evident. This time the mass on the left tail (worker only neighborhoods) is notable, although the mass on the right tail (resident only neighborhoods) is also evident. And again, the contrast with Taipei is clear.

Table 1 shows the summary statistic of the extent to which neighborhoods are specialized in each city. The first column shows the specialization in terms of land use. In each city, we calculate the specialization in terms of land use as the sum of the absolute value of

[11] When the index is positive, it is normalized by the ratio of the total residential space in the city to the total space in the neighborhood. When the index is negative, it is normalized by the ratio of the total commercial space in the city to the total space of the neighborhood.

Figure 4: Specialization of Neighborhoods: Taipei vs. US



Notes: Left panel shows the distribution of $H_{Ri}/H_R - H_{Ci}/H_C$ in each neighborhood. We normalize so that the maximum in each city is 1 (complete specialization in residential space) and the minimum is -1 (complete specialization in commercial space). Right panel shows the distribution of $L_{Ri}/L_R - L_{Ci}/L_C$ in each neighborhood. We normalize so that the maximum is 1 (neighborhood only has local residents) and minimum is -1 (neighborhood only has workers). U.S. is average of 36 American MSAs.

the index shown in the left panel of Figure 4 divided by 2. This statistic thus ranges from 0 (every neighborhood is mixed use in terms of land space) to 1 (every neighborhood is completely specialized in terms of space). The index is 0.65 for the average American city; the highest is 0.79 in Chicago and the lowest is 0.56 in Austin. The same index for Taipei is 0.49, which is lower than the American city that is the least specialized (Austin).

The second column in Table 1 shows the summary statistic of the specialization in terms of workers vs. residents, calculated analogously. This statistic is 0 when every neighborhood in the city has exactly the same number of workers and residents. It is 1 when every neighborhood either only has workers or only has residents. The specialization index in terms of workers vs. residents of the average American city is 0.62 and only 0.5 in Taipei. The most specialized American city is Las Vegas (0.69) and the least specialized is Portland (0.56). Note here that even the least specialized American city is more specialized compared to Taipei.

In the model, the distribution of space (residential vs. commercial) and people (workers vs. residents) between neighborhoods is driven by two forces: the dispersion

Table 1: Specialization in Space and People

	Residential vs. Commercial Space	Workers vs. Residents
Average U.S. City	0.65	0.62
Taipei	0.49	0.50
Lowest	0.56 (Austin)	0.56 (Portland)
Largest	0.79 (Chicago)	0.69 (Las Vegas)
Median	0.64 (Miami)	0.62 (San Antonio)

Note: Column 1 calculates the specialization of neighborhoods in commercial vs. residential space as $\frac{1}{2}\sum_i \left|\frac{H_{Ri}}{H_R} - \frac{H_{Ci}}{H_C}\right|$. Column 2 calculates the specialization of neighborhoods in people vs. workers as $\frac{1}{2}\sum_i \left|\frac{L_{Ri}}{L} - \frac{L_{Ei}}{L}\right|$.

of comparative advantage (A relative to B) between neighborhoods and residential vs. commercial zoning, as captured by the dispersion of τ. We now proceed to decompose these two forces.

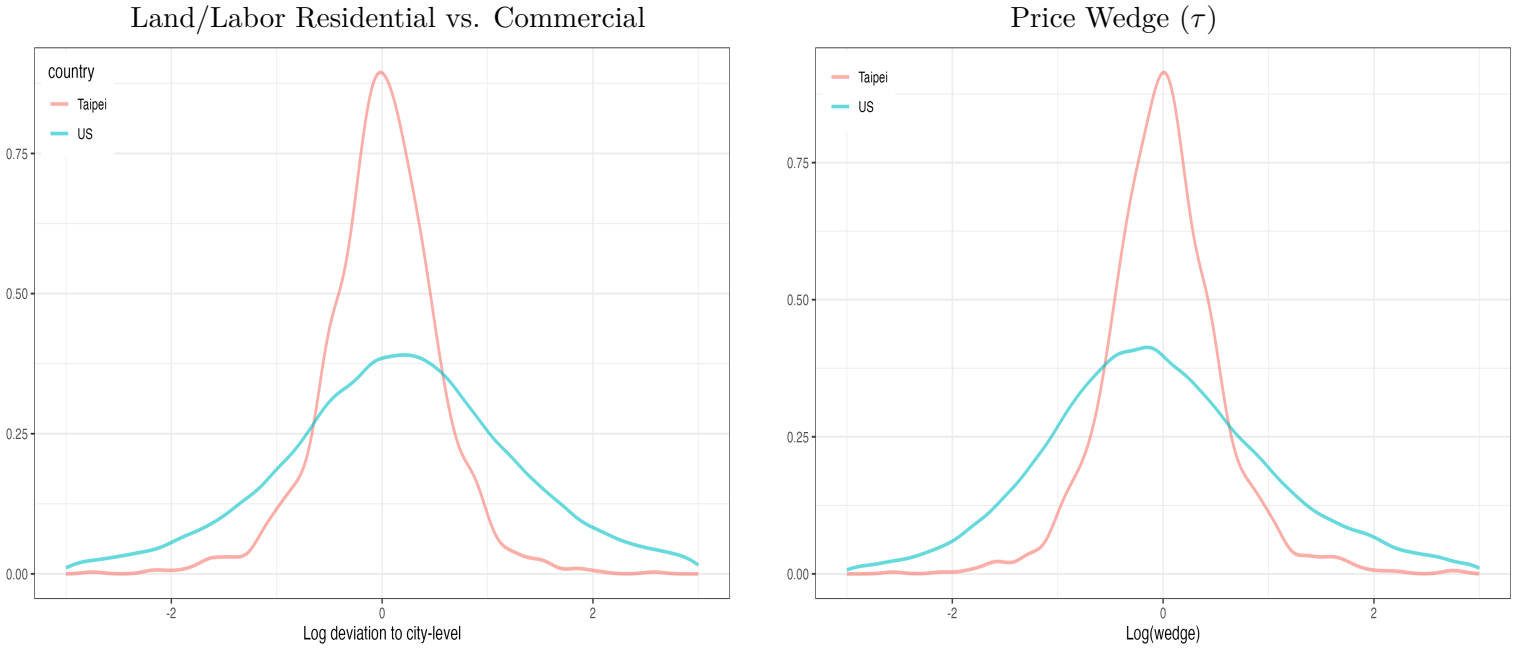
Remember that τ drives a wedge between the residential floor space per resident and commercial floor space per worker. We therefore start by showing the distribution of this gap between neighborhoods. Figure 5 (left panel) shows that the dispersion of this gap is significantly wider in American cities compared to Taipei.

The right panel in Figure 5 directly shows the dispersion of the price wedge τ implied by the data on floor space per person in the left panel. Following equation (4), we impute τ from the gap between residential floor space per resident and commercial floor space per worker, after adjusting for the ratio of local income to the local wage. The figure clearly shows that the dispersion of the price wedge between residential and commercial land is significantly larger in U.S. cities compared to Taipei.

Table 2 (first column) shows the 90-10 gap of the price wedge. In the average American city, the 90-10 gap of the wedge between commercial and residential space between neighborhoods is 2.74 log points. In Taipei the comparable number is 1.25 log points. The American city with the highest price dispersion is Detroit, and the city the lowest price dispersion is Austin. However, the price dispersion in Austin is still about 1 log points higher than in Taipei.

The second force that drives the observed specialization in land use and workers vs. residents is the dispersion of comparative advantage in production vs. residential amenities. Figure 6 shows the distribution of $\log A/B$ in Taipei and in the average American city. Two things are evident. First, the dispersion in $\log A/B$ is smaller in

Figure 5: Distribution of Land/Labor and Price Wedge (Residential vs. Commercial)



Notes: Left panel shows the distribution of log $\frac{H_{Ri}/L_{Ri}}{H_{Ci}/L_{Ei}}$ in each neighborhood. Right panel shows the implied distribution of log τ_i across neighborhoods imputed from equation (4). U.S. is average of 36 American MSAs.

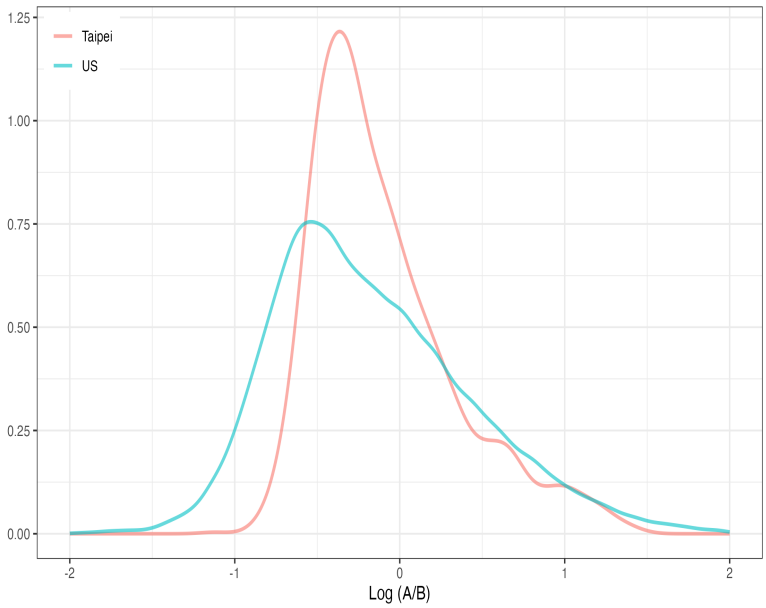
Table 2: 90-10 Gap in Local Price Wedge and Comparative Advantage

	Price Wedge (log τ)	Comparative Advantage (log A/B)
Average U.S. City	2.74	1.50
Taipei	1.25	1.11
Lowest	2.27 (Austin)	1.16 (Seattle)
Largest	3.29 (Detroit)	1.84 (Detroit)
Median	2.72 (LA)	1.52 (Orlando)

Note: Table shows the 90-10 Gap in log τ and log A/B across neighborhoods in each city. Average U.S. city is the average of 36 MSAs.

Taipei than in American cities. Second, the distribution is right-skewed both in Taipei and in the US.

Figure 6: Distribution of Comparative Advantage in Production vs. Amenities



Notes: Figure shows the distribution of $\log A/B$. U.S. is average of 36 large MSAs.

Table 2 shows the 90-10 gap of $\log A/B$. The 90-10 gap in $\log A/B$ is 1.5 in the average American city; the city with the largest dispersion is Detroit and the smallest is Seattle. The table also confirms the visual impression from Figure 6 that the dispersion of $\log A/B$ is smaller in Taipei compared to American cities.

We end this section with a decomposition of the relative importance of the price wedge τ vs. comparative advantage A/B on the specialization facts shown in Figure 4. We do this by regressing $\log \tau_i$ on the neighborhood specialization index on residential vs. commercial space shown in Figure 4 (left panel). The R^2 from this regression can be interpreted as the share of the variance of the specialization index due to τ_i (and where we equally split the covariance between τ_i and A/B). The R^2 from this regression averages 0.25, so the dispersion in τ_i "accounts" for about a quarter of the specialization in commercial vs. residential space. The R^2 from a similar regression in Taipei is only 0.09, so obviously much lower than in the average American city. The city with the highest R^2 from this regression is Denver (0.37) and the lowest is Austin (0.09).

6 Effect of Zoning

In this section, we use the estimates of $\log \tau$, $\log A$, and $\log B$ to ask the question: "what is the effect of the price wedge on aggregate welfare and on the specialization of neighborhoods in American cities?"

The problem with interpreting our estimates of τ as the price wedge is that τ can also reflect other forces. Specifically, the dispersion in τ can be driven by dispersion in the Cobb Douglas shares in production vs. utility (remember we assumed $\frac{\gamma\,\alpha}{1-\alpha}$ is constant across neighborhoods in a given city to infer τ). It can also reflect measurement error in the ratio of residential space to residents to commercial space in the number of workers.

This is where the estimates of τ dispersion in Taipei is important. Remember our evidence that there is effectively no commercial vs. residential zoning in Taipei. Therefore, the estimates of τ dispersion in Taipei can be interpreted as the effect of measurement error and model mispecification. We therefore estimate the effect of zoning in American cities as the answer to the following question: "what is the effect of changing the dispersion of τ in each American city so that the dispersion is equal to that of Taipei?" That is, in each American city, our counterfactual lowers the dispersion of τ not to zero but to that observed in Taipei.

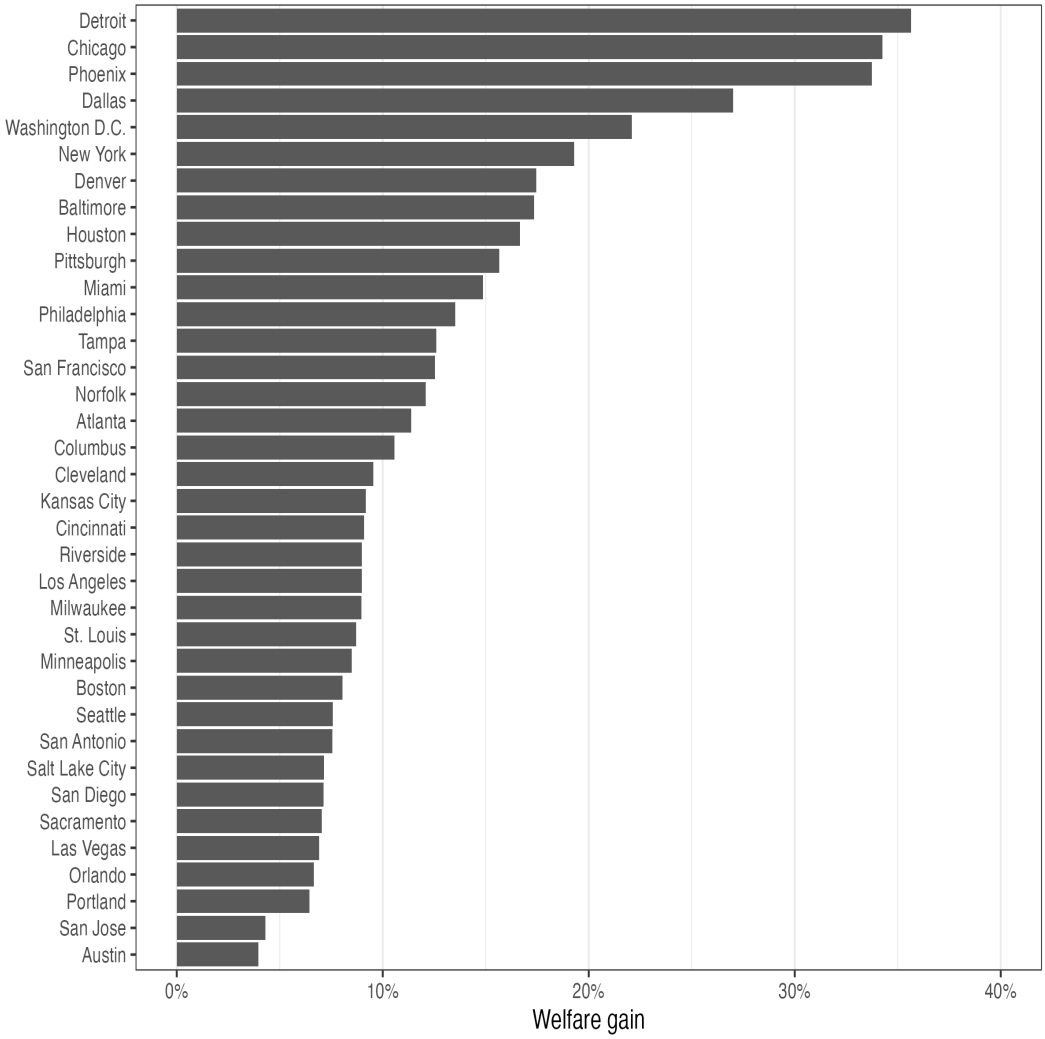
Figure 7 shows the results of this exercise for each American city. The welfare gains from eliminating residential vs. commercial zoning is 13.6% in the average American city. This number ranges widely between cities depending on the distribution of τ, A, and B across neighborhoods. The city with the largest welfare loss from zoning is Chicago (37% welfare loss). The welfare loss is still substantial in the city with the smallest welfare loss (Austin, with a 4% welfare loss).

Figure 8 shows the effect of eliminating zoning on the specialization of neighborhoods. The left panel shows the specialization of neighborhoods in the average American city in terms of residential vs. commercial space. For comparison, the figure replicates the actual distribution of the specialization of neighborhoods in terms of land use. The right panel shows the specialization of neighborhoods in terms of workers vs. residents. Again, the figure replicates the actual distribution of the share of workers vs. residents in each neighboorhood. Both figures show that the distribution under the no-zoning counterfactual becomes is less dispersed compared to what we see in the data.

Figure 9 shows the summary statistic of the specialization index under the "no-zoning" counterfactual, where the aggregate specialization index for each city is calculated as in Table 1 plotted against the specialization index measured in the data. So each observation in the figure is the summary statistic of specialization for one U.S. city. The left panel shows this number for specialization in terms of residential vs. commercial space. The right panel shows this number for specialization in terms of workers vs. residents.

The figure shows that in almost all U.S. cities in the left panel, the observations lie below the 45 degree lie. That is, there is less specialization in commercial vs. residential space under the no-zoning counterfactual compared to what is observed in the data. The pattern is less stark in the right panel, which shows the aggregate

Figure 7: Welfare Gain from Eliminating Residential vs. Commercial Zoning

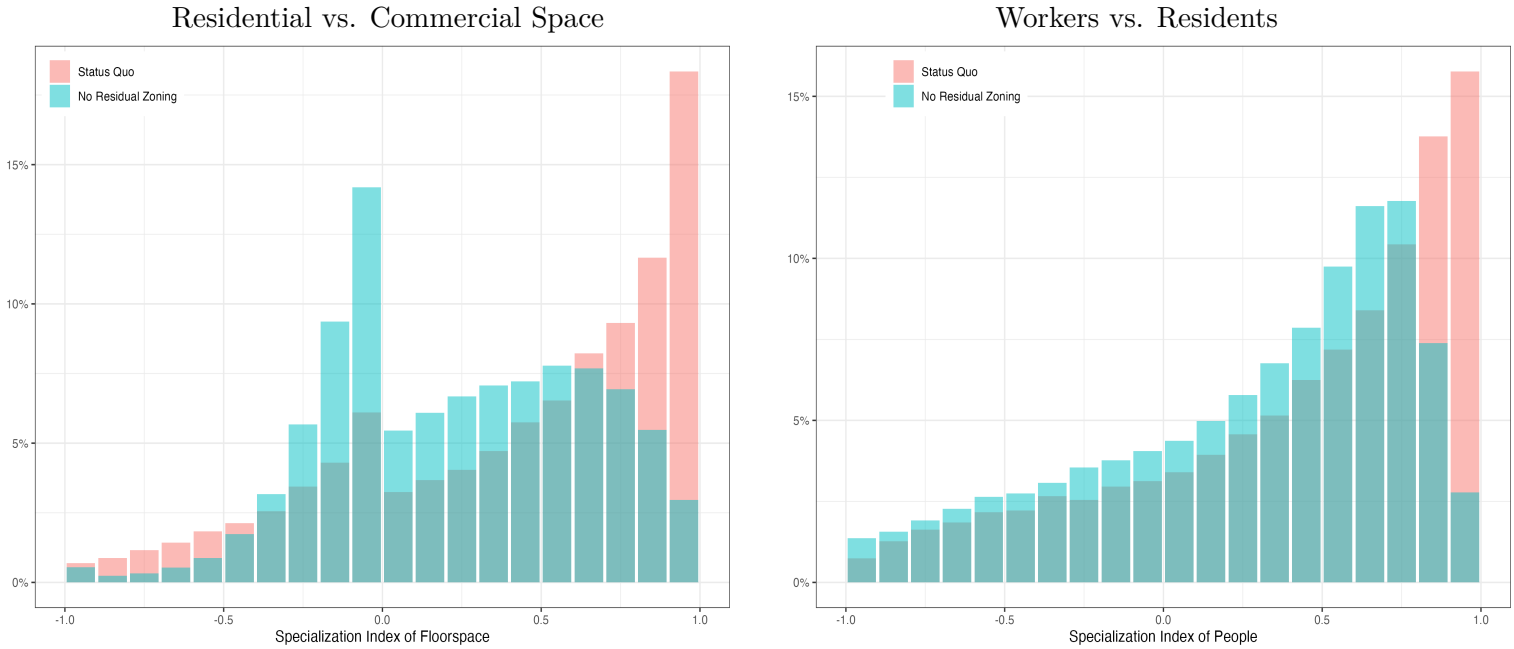


Notes: Figure shows the welfare gain from setting the dispersion of τ in each city equal to that of Taipei.

index of specialization in terms of workers vs. residents. The majority of U.S. cities lie below the 45 degree lie, indicating that the no-zoning counterfactual implies more mixed neighborhoods. However, there are a three cities – Chicago, Dallas, Detroit – that would be more segregated in terms of workers vs. residents under the no-zoning counterfactual. counterfactual.

We note that even under the no-zoning counterfactual, U.S. cities would still be more specialized compared to Taipei. This is because, as shown in Table 2, the dispersion of

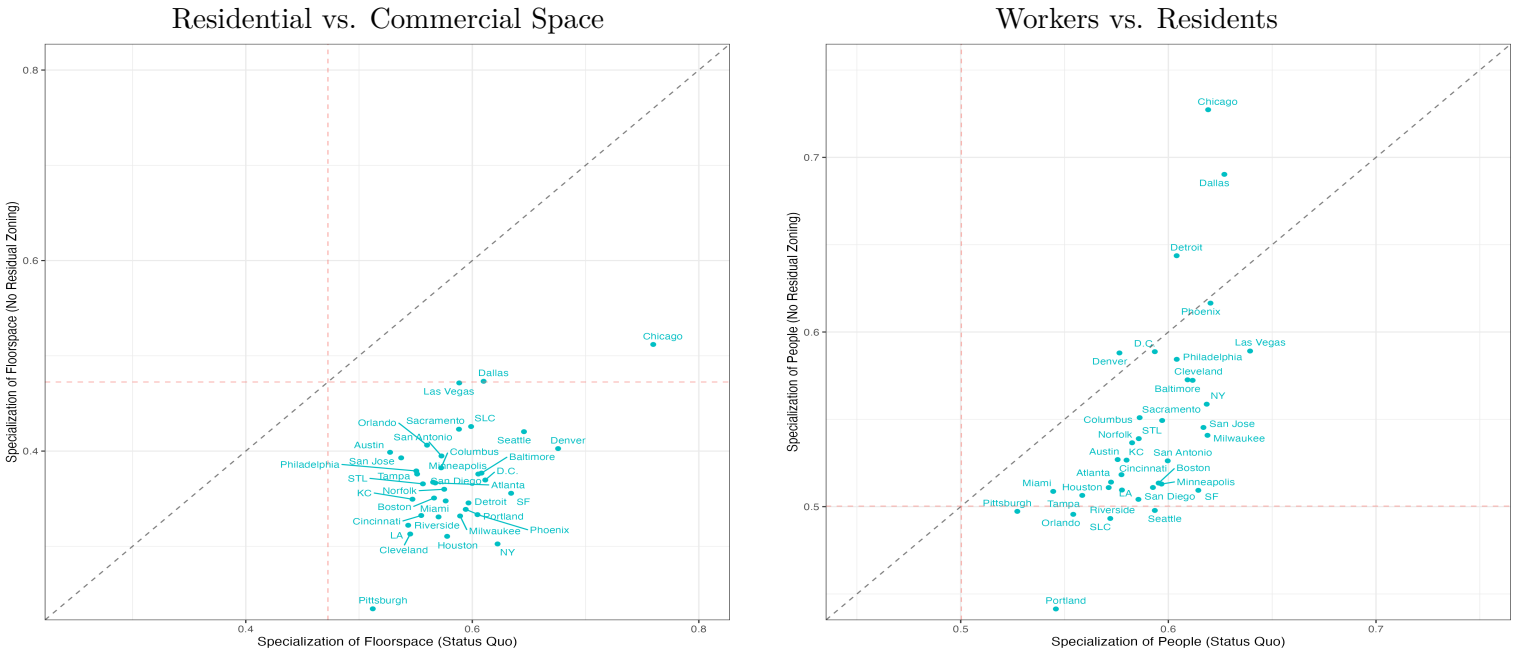
Figure 8: Specialization of Neighborhoods with No-Zoning



Notes: Figure shows the counterfactual specialization of neighborhoods when the dispersion of τ in each city is set to that in Taipei (see text for details). For comparison figure also shows the actual specialization of neighborhoods observed in the data. Left panel shows the distribution of $H_{Ri}/H_R - H_{Ci}/H_C$; right panel shows the distribution of $L_{Ri}/L_R - L_{Ci}/L_C$ in each neighborhood. Figure shows the average distribution across neighborhoods in 36 American MSAs.

comparative advantage in Taipei is smaller than in American cities. We have nothing to say about why the dispersion is smaller in Taipei, or why the dispersion differs between American cities for that matter. That is obviously an important question for future work.

Figure 9: Aggregate Index of Specialization with No-Zoning



Notes: Figure shows the aggregate index of specialization of each city under the no-zoning counterfactual (y-axis) vs. the actual index of specialization of each city (x-axis). Left panel calculates the index of each MSA as $\frac{1}{2}\sum_i |H_{Ri}/H_R - H_{Ci}/H_C|$; right panel shows $\frac{1}{2}\sum_i |L_{Ri}/L_R - L_{Ci}/L_C|$ in each MSA.