A Field Guide to Urban Economics

Matthew Turner

This version: July 31, 2025

Contents

Chapter 1

The Monocentric City Model

1.1 The Realtor's mantra and spatial equilibrium

The late Lord Harold Samuel, a real estate tycoon in mid-20th century Britain, is reported to have said: "There are three things that matter in property: location, location, location." [1] This sounds like a punchline, but it is important for two reasons. First, it is an expert's observation about how the world works. Second, it highlights what is special about the economics of cities. You can't study how cities are organized without thinking about where things are.

Consider two houses, house A and house B, alike in every detail except that the house B is downwind from a landfill and house A is not. Both are empty and available for rent, and there are many renters looking for houses in the neighborhood. The renters are all alike in the way they value both houses, and the landfill causes them all one dollar's worth of unhappiness. What should happen?

With many people looking for houses, both houses should end up rented. Moreover, the difference in their prices should be exactly equal to the one dollar of unhappiness that all renters suffer from being downwind from the landfill. This is exactly the Realtor's mantra. The difference in rent between the two houses is completely determined by their locations, downwind from a landfill and not.

This argument just tells us that the rent for house A is a dollar more than for house B. It doesn't tell us the actual level of the rent for either house. How do we set the level of prices? It must be that the households in the two houses don't want to move to wherever the large pool of unsuccessful renters landed, some offstage "outside option", so the rent of house A should reflect the benefit from living in house A relative to this outside option, with the rent for house B a dollar less.

[1] For more detail on the origins of this saying, see William Safire's June 26, 2009, column in the New York Times.

We have just worked out a simple model of spatial equilibrium for the housing market. In this equilibrium, identical people choose their favorite location from among the locations (and prices) available, and real estate prices adjust so that no one wants to move. An implication of this sort of equilibrium is that real estate prices are determined by the value of differences in the services that each location can provide to its occupant. That is, "location, location, location." As a starting point, this looks pretty good. The implications of our little economic model line up with the way the professionals think real estate markets should behave.

This model of how real estate markets work has an important implication for welfare. Which of the two successful renters is better off, the one downwind from the landfill and paying one dollar less rent, or the one without the nearby landfill? Suppose that the household near the landfill is worse off than the other. Recalling that the households are the same to start with, then this household has made a mistake. They should have offered the landlord of the first house an extra penny and moved into the house that is not near the landfill. At market rents, it must be that the two households are indifferent between the two houses. Lower rent must compensate the household downwind from the landfill for the unhappiness this causes or this household will move away. Even though the landfill affects just one of the two households, both are equally well off.

Suppose that, to promote environmental justice, you are asked to vote for a ballot initiative that will clean up the dump so that the downwind household can no longer notice it. If this ballot initiative is passed, who benefits from the cleanup? Absent the landfill, the resident of house B has an identical house to house A, but pays 1 dollar less in rent. This means that the resident of house A,or one of the many people who did not find a place to live in the neighborhood, should bid up the price of house B to equal house A. This may take a while, the lease contract may run for a year, it may take people a while to figure out that the dump is cleaned up, but in a perfect world, this would happen pretty fast.

So who benefits from the dump clean up? Is it the household downwind from the landfill? No. This household is indifferent between the two houses before the clean-up. After the clean-up, and rent adjustment, it is still indifferent between the two houses. The winner from the landfill clean-up is the landlord of the downwind house. Will this change your vote?

Now consider two more examples. Suppose that instead of being downwind from a landfill, the second house is next to a gangster who collects one dollar in protection money every month, but is otherwise pleasant and unthreatening. This should operate much like being downwind from a landfill. The second house comes with a one dollar monthly cost and the first does not, so the rent for the second house must be one dollar less. The real estate market should deal with both noxious neighbors in exactly the same way.

Now for the interesting case. Suppose that the second house comes with a one dollar per month property tax bill, payable by the tenant, and the first does not. This has exactly the same implications for the resident of the second house as does the gangster neighbor; one dollar out of pocket each month. It should therefore have the same implications for the real estate market. That is, the rent in the second house should be exactly one dollar lower than the first, and the two households should be indifferent between the two locations.

Why is this interesting? It means that the property tax (1) does not affect the welfare of the people who live in taxed houses, and (2) that the property tax does not affect the total cost, rent paid to the landlord plus tax paid to the government, of a property. This is not how taxes on most other goods work, they usually raise prices at least a little.

This leads us to two bits of jargon. The first is easy. In all three examples, landfill, gangster, and property tax, we say that real estate prices "capitalize" whatever is special about the second house. That is, prices adjust to reflect differences in the value of the services provided by each location.

The second is "economic rent", and it is a little slippery. Returning to the landfill example, we recall the household that occupies upwind house A is willing to pay an extra dollar of rent to avoid downwind house B. The need for jargon arises when we replace the landfill with the gangster or tax payment. Here, the payment the tenant in house B makes to the landlord is one dollar less than that of the tenant in house A, but the *total* payments for house A and B are the same. Without the landfill, the value that a tenant gets from each house is the same. The difference is that the landlord for house A receives the money equivalent of their tenant's value of living in the house, but the landlord for house B splits this value with the city government or the gangster. "Economic rent" is the value that a household gets from living in house A or house B. It is sometimes different than "contract rent", what the tenant pays the landlord. In these examples, contract rent and economic rent coincide for house A, but economic rent is divided between a contract rent payment and a payment to the city government or the gangster in house B. From here on, "rent" always means "economic rent", unless I explicitly note otherwise.

The object of this book is to understand the way people make the decisions that build and organize the cities where most of us live. The rest of this Chapter, and much of the rest of the book, revolves around applying the notion of spatial equilibrium that we have worked out here. That is, we ask what happens when people choose their favorite location from among the locations available and real estate prices adjust so that no one wants to move. However, instead of considering houses that are different from each other because of their proximity to gangsters and landfills, we consider houses that differ in their proximity to the center of a city where people work. This will give rise to one of the main theoretical tools that we have for thinking about the economic geography of cities, the "monocentric city model".

1.2 Land rent gradients in real life

If we study the allocation of sugar donuts, we need just one price, the price of a sugar donut. But if we are studying the the price of land, we need a price for each location. If we think space is continuous, then we need a continuum of prices. A little more formally, in most of the rest of economics, prices are scalars. In urban economics, they are functions. We don't have a price of land, we have a land price function. This function assigns a price to each location. This land price function is often called a "land price gradient".

The goal for this book is to learn about the economics of cities, and land markets are central to this project. Figure 1.1 starts us off by describing land price gradients in two Japanese and two French cities. All four panels of figure 1.1 show how land prices change with distance to the city center. The two top panels show how land prices fall with distance from the city center for two cities in Japan in 1991, Hiratsuka and Yokohama. They fall fast. In panel (a) we see that a square meter of land near the center of Hiratsuka sells for 2 to 3 million Yen. A mile away, this price falls to half a million, and by 8 miles away, it has fallen to 100,000 or less. Panel (b) shows similar data for Yokohama, a much bigger city. Here, land near the city center sells for 10-20 million Yen per square meter, and shows the same rapid decline with distance to the center. Notice the mismatched units on both figures, metric on the y-axis and imperial on the x-axis.

The bottom two panels of figure 1.1 differ from the top two in two ways. First, they are describing cities in France in 2012 instead of Japan in 1991. Second, both axes are in logarithms rather than levels. Panel (c) plots the logarithm of land prices in Paris against the logarithm of distance to the center. Panel (d) is the same, but for Dijon. These two figures also show a clear decline in land prices with distance to the center, but because the data is presented as logarithms, it is hard to tell if the decline is as fast as it is in the Japanese cities. For this, we need to do a little math.

Let R indicate the price of land and x be the radial distance from the center of the city. The Japanese figures plot R against x. The French figures plot $\ln R$ against $\ln x$. To compare these two types of plots, you need to remember the rules for logarithms.

$$\ln R = A + B \ln x \tag{1.1}$$

$$\implies \ln R = \ln e^A + \ln x^B \tag{1.2}$$

$$\implies \ln R = \ln e^A x^B \tag{1.3}$$

$$\implies R = e^A x^B \tag{1.4}$$

Figure 1.1: The relationship between land prices and distance to the center in four cities

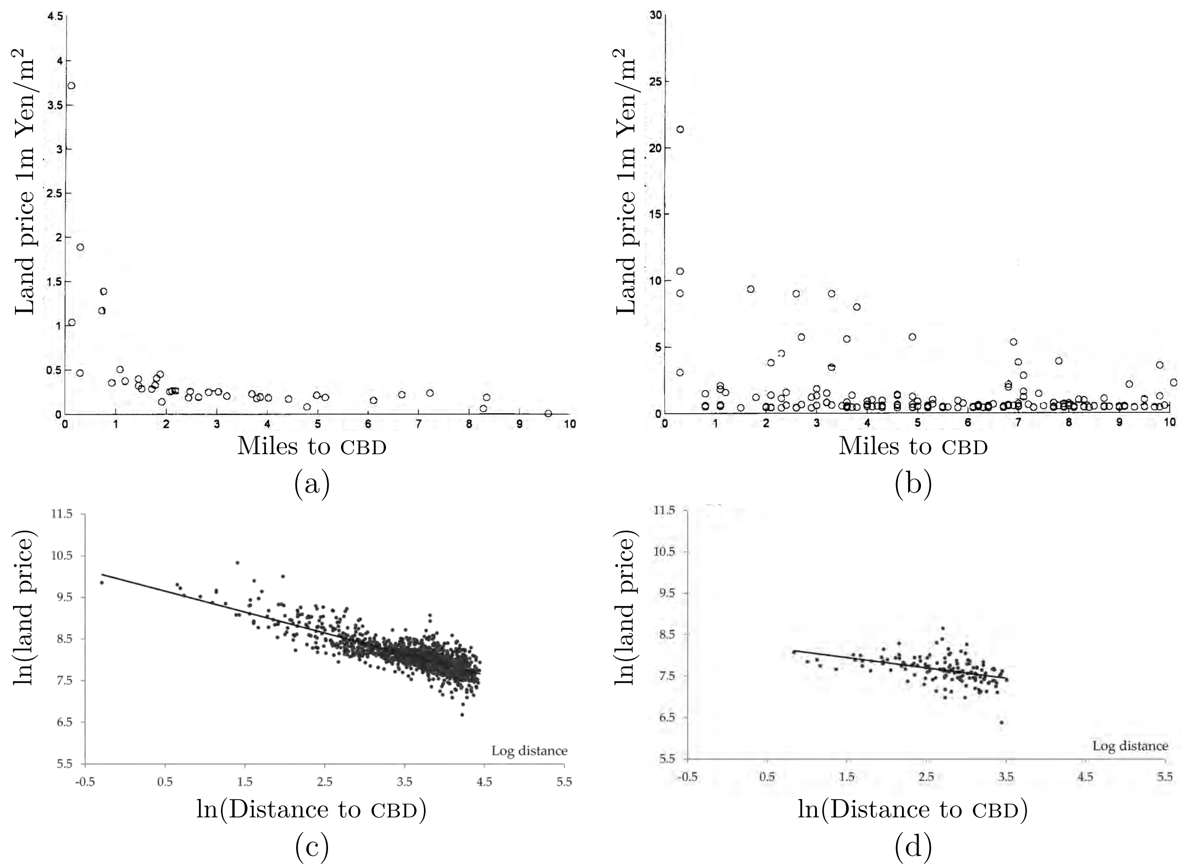

Note: *(a) 1991 land prices in Hiratsuka, Japan; (b) 1991 land prices in Yokohama, Japan; (c) logarithm of 2012 land prices in Paris, France; (d) logarithm of 2012 land prices in Dijon, France. Figures (a,b) from Lucas [2001] show how land rent declines (very fast) with radial distance from the center of two Japanese cities. Figures (c,d) from Combes et al. [2019] show the decline of the natural logarithm of rent with the logarithm of radial distance to the center.*

Equation (1.1) describes the line plotted for the two French cities, the logarithm of land price against the logarithm of distance to the center. Equation (1.2) uses the fact that logarithms and exponentiation are inverses, that is, $\ln e^x = x$. Equation (1.3) uses a rule of logarithms, $\ln(x) + \ln(y) = \ln(xy)$. The last equation uses the fact that logarithms and exponentiation are inverses again. This last equation is the one that is plotted for the Japanese cities, so we see that the the two different looking pairs of graphs are actually plotting the same information, but represented differently.

Eye-balling the figure for Paris, we see that the intercept is about 10 and the slope

Figure 1.2: Comparing plots in logarithms and levels

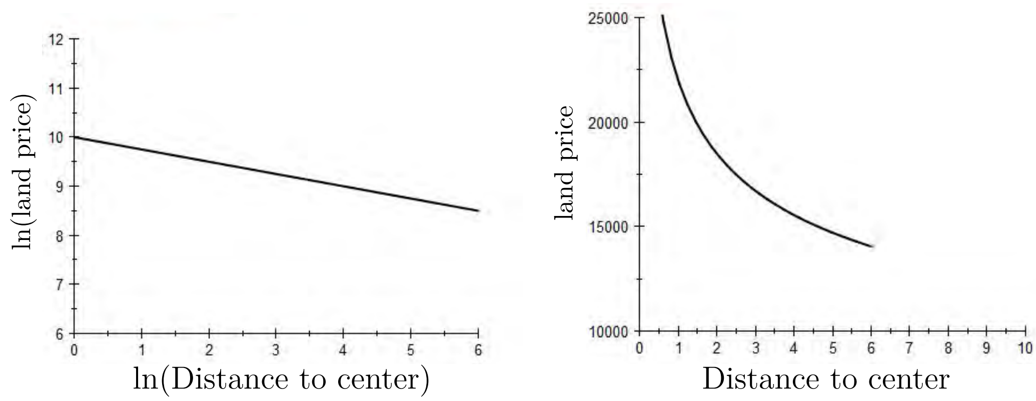

Note: *The left panel plots equation* (1.5)*. The right panel plots equation* (1.6)*. Both panels describe the same relationship between R and x, but the one on the left is in logarithms and the one on the right is in levels.*

about $-\frac{1}{4}$. Writing this out, we have

$$\ln R = 10 - \frac{1}{4} \ln x. \tag{1.5}$$

We can rearrange this to get

$$R = e^{10} x^{-\frac{1}{4}} \approx 22,000 x^{-\frac{1}{4}}. \tag{1.6}$$

The left panel of figure 1.2 plots the first of these equations. The right panel plots the second. Once we convert the French data from logarithms to levels, we see the same rapid radial decline in land prices that we see for Japanese cities.

So, land rent behaves the same way in France as it does in Japan. This is pretty neat. It did not have to be true. In fact, cities almost everywhere show this sort of log-linear decline in rent with distance to the center.

Now, two asides. First, economists often find the world is well described by log-linear relationships like the ones illustrated for the French cities in figure 1.2, so it's worth learning how to go back and forth between logarithms and levels (as we've just done). Second, log-linear relationships have another advantage. The coefficient B on $\ln x$ in equation (1.1) is an elasticity. It tells us the percentage change in rent that results from a one percent change in distance (see box 6.7.1). Elasticities are handy because you don't need to keep track of the units that you use to measure x and R, or whatever variables you are interested in. You can use meters to describe your y-axis and miles for the x-axis and, if you plot your data in logarithms, you won't get caught.

Box 1.2.1: Regression coefficients in log-linear specifications are elasticities

Suppose that we can write n as a log-linear function of r. That is,

$$\ln n = A + B \ln r. \tag{1.7}$$

To see that B is an elasticity, suppose we increase r by 1%, from r^0 to $r^1 = 1.01r^0$, all else equal. Then, we have

$$\ln n^0 = A + B \ln r^0 \tag{1.8}$$

and

$$\begin{aligned} \ln n^1 &= A + B \ln r^1 \\ &= A + B \ln(1.01) r^0 \\ &= A + B(\ln(1.01) + \ln r^0). \end{aligned} \tag{1.9}$$

Subtracting equation (1.8) from (1.9), we have

$$\ln n^1 - \ln n^0 = B \ln(1.01).$$

or

$$\ln \frac{n^1}{n^0} = B \ln(1.01).$$

Next, define ρ as the proportional change in n, so that $\frac{n^1}{n^0} = 1 + \rho$, and recall that $\ln(1 + x) \approx x$ for x small, and we have

$$\rho = B \times 0.01.$$

Multiplying by 100, we have that $100\rho = B$. That is, B is the percentage change in r that results from a 1% change in n. In the jargon, B is the elasticity of n with respect to r.

1.3 The monocentric city model

We have so far established two ideas. First, that the rental prices of property ought to adjust to reflect differences in the value of living at the property in such a way that no one wants to move. Second, the price of land falls rapidly as we move away from the center of cities.

The monocentric city model explains the second fact as a consequence of the first. That is, it assumes the price of real estate changes in order to keep people indifferent between all available locations, just as in our example with house A and house B and uses this assumption to explain the decrease in real estate prices as we get further from the center. With our example of the landfill still in mind, you can guess how this is going to work. For prices to fall with distance from the center, something about more remote locations has to get worse. In the monocentric city model, the thing that gets worse with distance to the center is the cost of commuting to work. This is the central intuition of the model; land prices fall with remoteness from the center to exactly compensate for a more costly commute.

Before we develop the model, two comments. First, the data on land prices presented in section 1.2 describe the price of land, how much you have to pay to obtain the services of the land forever. In our examples, we've considered the rental price of land, what you have to pay to obtain the services of the land for some definite fixed time. For now, let's just name these two Asset Prices and Rental Prices, respectively, and note that although they are not the same, they are close relatives. We'll work out the relationship between them later, but for now, you can think of them as synonyms.

Second, in order to talk about prices (or rents) declining with distance from the center, we need to locate the center. It turns out that this is a surprisingly well defined concept. Ask yourself, and two or three other people, where the center of your hometown is. You will almost surely get the same answer from everyone. In Providence, where I am writing this Chapter, it is the plaza across from city hall. We will refer to this central location and the area around it as the "Central Business District", or CBD.

Imagine a city located on a featureless plane or along a line. We begin with a linear city because it is a little simpler. Indicate locations on the line with x. There is a CBD located at $x = 0$ and $|x|$ is distance to the CBD, with $x < 0$ for a location to the left of the CBD and conversely. There is one unit of land at each x.[2] The city is populated by identical households (or workers), all of whom commute to the CBD where they earn wage w. Commuting costs t per unit distance, and so a household living at x must pay $2t|x|$ in commuting costs to the CBD and back. All households occupy a parcel of fixed size, $\overline{\ell}$, at the location x that they choose. Households use their wage to pay the land rent for their parcel, $R(x)\overline{\ell}$, to purchase a composite consumption good, c, and to pay the cost of commuting, $2t|x|$. Households derive utility from the consumption good according to the utility function $u(c)$, and we require that $u' > 0$.

Land not used for urban residences remains in agricultural use where a farmer is

[2] This is a little bit fishy but let us avoid some arcane math. How can there be one unit of land at a point on a line? Strictly, and if you know some probability theory, there is a uniform "density" of land.

always willing to pay a reservation land rent of $\overline{R}$, the "agricultural land rent". The total population of the city is N and the price of the consumption good, c, is set to $p_c = 1$ (so we don't need to keep track of it). Finally, all land rent is collected by "absentee landlords". This is an important bit of fiction. It means that all land rent leaves the model and we don't need worry about the messy problem of keeping track of who gets to spend it.

We make two assumptions about how households behave. First, that they choose their location x so as to make themselves as well off as possible. That is, households solve,

$$v(c,\, x) = \max_{c,\, x} u(c) \qquad \text{s.t. } w = c + R(x)\overline{\ell} + 2t|x|. \tag{1.10}$$

This says that households choose their residential location x to maximize their consumption, subject to their budget. This is the standard assumption about rationality in economics: People make themselves as well off as they can given the choices available to them. As long as u is an increasing function this maximization problem is trivial; use everything left from the wage after paying for rent and commuting to buy the composite consumption good. No calculus required.

The second assumption we make is that no one wants to move. Formally, we require that households get the same utility at every x, and call this utility level $\overline{u}$. That is,

$$v(c,\, x) = \overline{u} \text{ at all occupied } x. \tag{1.11}$$

This is sometimes called a "free mobility condition". If it is free to move, we expect people to move for any tiny improvement in their utility, and so utility must be the same everywhere. We often call $\overline{u}$ a "reservation utility level".

This is the complete statement of the model.

Before we work out the implications of these assumptions, three comments are in order. First, as anyone who has ever moved knows, moving is not free, and so starting from an assumption of "free-mobility" seems inauspicious. But the assumption is better than it seems. Consider someone who has already decided to move to a new city. All of their stuff is in the mail or on a truck. For this person, choosing a house on one block or another is essentially free. Moreover, this is one of the people who is actually participating in the market and helping to set prices. This is what the free mobility assumption really requires. Not that everyone can move costlessly, but that the subset of households who are buying or renting properties can do so. This seems much easier to defend.

Second, the monocentric city model comes in two varieties, "open city" and "closed city". In an open city model, city population adjusts until the free mobility condition is satisfied and everyone in the city is indifferent between all locations in the city and the outside option. In a closed city model, we fix the population of the city and let the utility level, $\overline{u}$, adjust so that households are indifferent between all locations in the city. These are stylized cases, and reality probably lies somewhere between.

Finally, according to data assembled by the United Nations Population Division, the share of the world's population that lives in cities from 1960 until 2020 rose from about 34% to about 56%, even as the world's population is increasing. As a starting point, we probably want to think about a model where the population of a city can adjust rather than one where it cannot. Open city models are also a little easier to work with, and so we'll start with this case.

We now turn to working out the implications of the monocentric city model. We would like to see what it implies about the extent of the city, its population, the land rent gradient, and the welfare of its residents.

To begin, invert the free mobility condition, equation (1.11), to find the level of consumption that households require to reach the reservation level of utility,

$$c^* = u^{-1}(\overline{u}). \tag{1.12}$$

For example, if $u(c) = c^{1/2}$ and $\overline{u} = 2$, then $(c^*)^{1/2} = 2$, or $c^* = 2^2 = 4$.

In a spatial equilibrium, everyone gets the same utility at every location, so consumption must be the same everywhere. Therefore,

$$w - c^* = R(x)\overline{\ell} + 2t|x|, \tag{1.13}$$

for all locations x. With wages and consumption fixed for all households, commuting costs and land rent must vary in such a way that they always sum to a constant. Implicitly, we're also assuming $w > c^*$. Otherwise, no one would live in the city at all.

We can use equation (1.13) to find the extent of an equilibrium city. The limits of the city are defined by the most remote points where a city resident values the land more highly than a farmer. That is, where a city resident is just willing to pay the reservation land rent $\overline{R}$. Let $\overline{x}$ denote the distance of these boundary points from $x = 0$. At this location, we must have

$$w - c^* = \overline{R\ell} + 2t|\overline{x}|.$$

At the edge of the city, the cost to commute is such that a household can just pay the reservation land rent and commuting costs, and still buy the reservation consumption bundle. Reorganizing, we have that

$$\overline{x} = \frac{w - c^* - \overline{R\ell}}{2t},$$

is the most remote occupied point to the right of $x = 0$, and the city extends from $-\overline{x}$ to $+\overline{x}$.

We see here how the open city assumption works. A household must obtain a utility level of at least $\overline{u}$ to live in the city. This means consumption of at least c^*. The price of land at the unoccupied location nearest the CBD must be $\overline{R}$, so the price of the best unoccupied parcel is $\overline{R\ell}$. This means that the most remote occupied location is one where a household can just afford c^* after bidding land away from a farmer and paying for their commute. In this sense, the city is "open", its size is determined by the number of people who choose to live there.

Because the city extends from $-\overline{x}$ to $\overline{x}$ and each household consumes an exogenously fixed amount of land, the population of the city is,

$$N^* = \frac{2\overline{x}}{\overline{\ell}}.$$

Note that we are here using the assumption that there is one unit of land at each x.

Using the equilibrium budget constraint, equation (1.13), and the equilibrium extent of the city, we can solve for the equilibrium rent gradient,

$$R^*(x) = \begin{cases} \frac{w-c^*-2t|x|}{\overline{\ell}} & \text{if} \quad |x| \leq \overline{x} \\ \overline{R} & \text{if} \quad |x| > \overline{x}. \end{cases} \tag{1.14}$$

If we restrict attention to locations in urban use, then this is just

$$R^*(x) = \frac{w - c^* - 2t|x|}{\overline{\ell}}. \tag{1.15}$$

This is a bit simpler to write out, and I'll often cheat and write it this way.

This finishes the solution of the model. Box 1.3.1 works out an example. Assuming that households optimize and that no one wants to move, and given a reservation utility level, a price of commuting, and a wage for work in the central location, we've derived the size of the city and its configuration, along with the land rent gradient. Our next step is to present the same argument graphically. This is a little easier and makes the intuition clearer. After that we want to think about some extensions of the model that make it a little more realistic, and to consider its other implications. So far, the model is able to predict the downward sloping rent gradient we observe. we'd like to have more predictions to check.

1.3.1 The monocentric city model in two pictures

The monocentric city model has a tidy graphical representation, given in figure 1.3. Let the x-axis indicate displacement away from the CBD at $x = 0$, and let the y-axis

Box 1.3.1: Example: Monocentric city

Suppose that, $u(c) = \ln(c)$, $\overline{R} = 0$, $\overline{u} = 0$ and $\overline{\ell} = 1$. Then, the household's problem is to choose location and consumption to solve,

$$\max_{c,\,x} \ \ln(c)$$
$$\text{s.t. } w = c + R(x) + 2t|x|.$$

Suppose the city is open, so that people migrate in or out until the utility level at all locations in the city is equal to the reservation utility level. Then, spatial equilibrium requires $\ln(c) = \overline{u} = 0$ so that $c^* = 1$ everywhere.
Using $c^* = 1$ in the budget constraint, we have

$$w - 1 = R(x) + 2t|x|$$

which means that,

$$R(x) = \begin{cases} w - 1 - 2t|x| & \text{if} \quad |x| \leq (w-1)/2t \\ 0 & \text{if} \quad |x| > (w-1)/2t. \end{cases}$$

The edges of the city are at $\overline{x} = \pm(w-1)/2t$ and because $\overline{\ell} = 1$ this means that $N^* = (w-1)/t$.

indicate units of consumption. Because w, c^*, $\overline{R}$ and $\overline{\ell}$ are the same at all locations, we can draw three horizontal lines, the first for the wage, at height w, the second for the wage net of consumption, at $w - c^*$, and the third at the value of land in agriculture, $\overline{R\ell}$.

Looking again at the budget constraint, we have

$$w - c^* = R(x)\overline{\ell} + 2t|x|.$$

That is, every household gets the same wage and enjoys the same level of consumption. Once they have paid for this consumption, the rest of their earnings go to land rent for their parcel of size $\overline{\ell}$ and to the cost of commuting.

A household living right at the the CBD, at $x = 0$, doesn't commute and so has zero commute expenditure. In order for this household to satisfy their budget constraint, the rest of their earnings, $w - c^*$, goes to land rent. It follows that $R(0) = (w - c^*)/\overline{\ell}$ and that the $x = 0$ household's total expenditure on land is $R(0)\overline{\ell} = (w - c^*)$. As we move away from $x = 0$, commute costs increase linearly, at a rate of $2t$ per unit distance. This gives us the dashed gray commute cost gradient. Expenditure on land

Figure 1.3: The monocentric city model in one easy picture

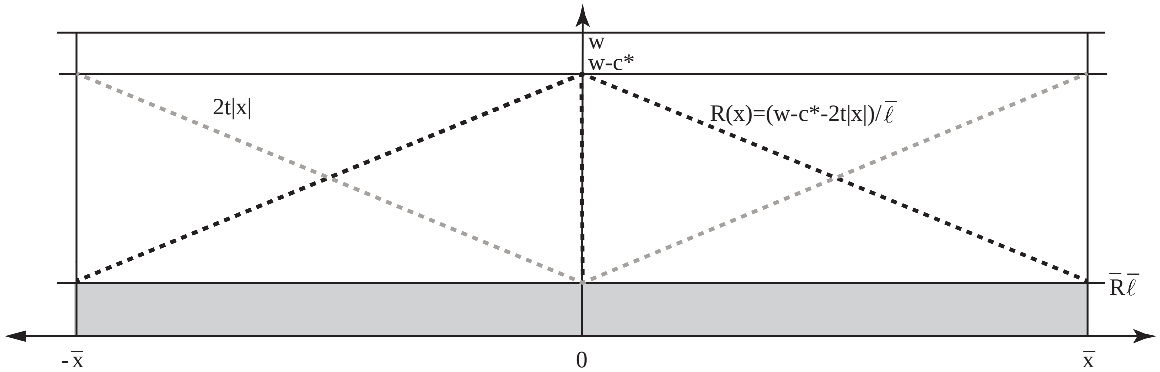

Note: *Illustration of the monocentric city model. x-axis is displacement from the* CBD *at $x = 0$ and y-axis is units of consumption. The wage, value of land in agriculture and household consumption level are all constant across all locations x. The gray dashed lines show how commute costs increase with distance from the* CBD *and the black dashed lines show how land rent decreases. The edges of the city occur where urban residents can no longer afford to commute and outbid farmers for land.*

decreases by an exactly offsetting amount. This gives us the the dashed black land rent gradient. The edge of the city occurs at a distance $\bar{x}$ from the CBD where, once a household pays for their commute, they have just enough left over to bid a parcel away from the farmers.

Early in the industrial revolution many cities were "mill towns". There was some big concentration of employment at the center, a mill, a collection of mills, a port or a railway depot. All the workers lived nearby and walked or took the train back and forth to the center. This is just the situation that the monocentric city model is meant to describe, and figure 1.4 shows that this is just how 19th century Providence was organized. Employment was highly centralized in the center, and there was no way to get back and forth to the CBD except for on foot or by train.

1.3.2 Three extensions

We now consider three extensions to the basic model. First, we consider a closed city equilibrium. Although we will work primarily with the open city model, most current research is based on models of closed cities, so this is an important case to work out. Second, we consider a circular city. The linear city assumption is obviously silly. We want to work out the circular city model to demonstrate that the extra realism doesn't actually change anything beyond making our math a little messier. Finally, we want to introduce the idea of "amenities" to our model. Amenities, here

Figure 1.4: The geography of Providence around 1896.

Note: *View of the city of Providence as seen from the dome of the new State House. Drawn by M. D. Mason, published in the Providence Sunday Journal, Nov. 15, 1896.*

some feature of the city that affects utility directly, like sunshine or pollution, are important for determining city size in reality and will play an important role in much of what we talk about later in the book.

Closed city equilibrium

The “closed city” version of the monocentric city model is exactly the same as the “open city” version we have worked out, except that instead of knowing the value for the reservation utility, $\overline{u}$, we know the number of people who live in the city, N.

Given population size N, because everyone consumes $\overline{\ell}$ of land, the length of the city must be $N\overline{\ell}$, and so the edges of the city must be at $|\overline{x}| = N\overline{\ell}/2$.

Once we know the most remote occupied locations, we can figure out consumption by requiring that rent at the edge of the city equal agricultural rent,

$$R(\overline{x}) = \frac{w - c^* - 2t|\overline{x}|}{\overline{\ell}} = \overline{R}.$$

Rearranging and substituting for $\overline{x}$, we have

$$c^* = w - \overline{R\ell} - tN\overline{\ell}.$$

In equilibrium, consumption must be the same at all occupied locations. That is,

$$c^* = w - R(x)\overline{\ell} - 2t|x|.$$

Equating these two expression for c^* and solving for $R(x)$, we have

$$R(x) = \frac{(\overline{R} + tN)\overline{\ell} - 2t|x|}{\overline{\ell}}.$$

This is complicated looking, but the intuition is the same as what we have already done. Rent adjusts so that income net of commuting and rent is the same everywhere. The difference is that in an open city equilibrium, the reservation utility is exogenous, and the population of the city adjusts until the marginal household pays just the agricultural rent. With a closed city equilibrium, population is fixed and utility adjusts.

Circular city

Suppose we relax the (silly) assumption that the city is on a line, and think about a circular city that is symmetric around a central point, still on a featureless plane, keeping everything else the same.

This barely changes the household's problem at all. We still have

$$\max_{c,\,x} u(c)$$
$$\text{s.t. } w = c + R(x)\overline{\ell} + 2tx$$

and, in an open city, $u(c^*) = \overline{u}$.

Although this problem looks the same as the one for the linear city. It is slightly different. In this case, x is radial distance to the center, in whatever direction, and can only be positive. In contrast, for a linear city, x and $-x$ refer to particular coordinates on the line. Practically, this means that we don't need the absolute value on x to state the circular city problem, and we need to keep in mind that x refers to all of the locations at distance x from the CBD, rather than to a particular point.

Consumption must still be the same everywhere in a spatial equilibrium,

$$w - c^* = R(x)\overline{\ell} + 2tx.$$

Let $\overline{x}$ denote the distance from the origin to the most remote occupied location. At this this distance from the CBD, we must have

$$w - c^* = \overline{R\ell} + 2t\overline{x}.$$

Reorganizing, we have

$$\overline{x} = \frac{w - c^* - \overline{R\ell}}{2t}.$$

This is the same as for the linear city, but it is now on the edge of a circular city.

The area of our circular city is $\pi\overline{x}^2$, so population is

$$N^* = \frac{\pi\overline{x}^2}{\overline{\ell}}.$$

This is the big (and completely obvious) difference between the linear and circular city. With the linear city, the extent of the city increases at the same rate as the population. With a circular city, area increases much faster than radius, so for a given increase in the radius of the city, a circular city accommodates a lot more people than does a linear city.

Amenities

Suppose our city has an amenity A that affects the utility of residents. This could be something like good or bad weather, crime, pollution, or parks. How does this affect equilibrium?

To illustrate ideas as simply as possible, suppose a household's utility is $u(Ac)$, almost just as before. So, $A > 1$ is something good, $A < 1$ is something bad. How does this change the open city equilibrium? With an open city, we have

$$u(Ac^{**}) = \overline{u}.$$

Reorganizing, we have,

$$c^{**} = \frac{1}{A}u^{-1}(\overline{u})$$

If $A = 1$ we get back the basic case we've already covered and $c^{**} = c^*$. If $A > 1$, then $c^{**} < c^*$. That is, if a city has an amenity that contributes to utility then

Figure 1.5: The monocentric city with amenities

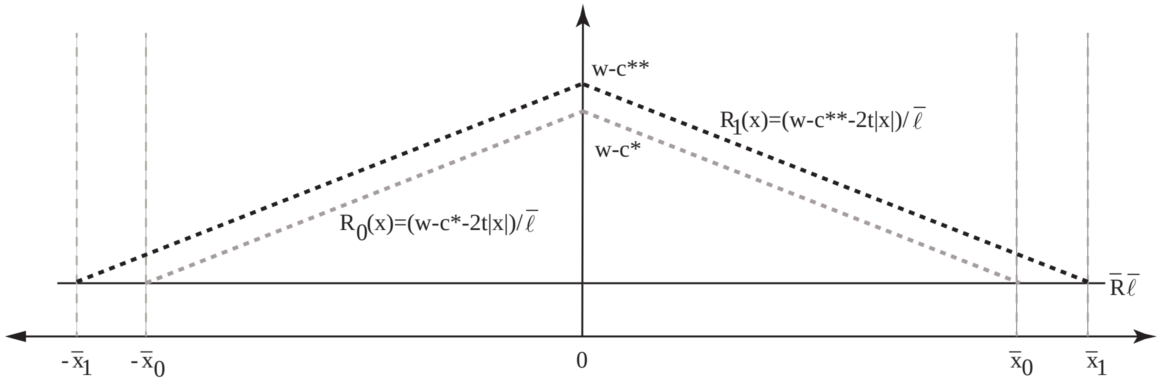

Note: *The dashed gray line illustrates the land rent gradient for the baseline case when there are no amenities, really when $A = 1$. The dashed black line illustrates the land rent gradient when the city has some amenity that complements ordinary consumption, $A > 1$. In the city with the amenity, rent goes up everywhere and the city expands. With the amenity, it is possible to hit the reservation utility level with a little less consumption. This leaves more income to be divided between land rent and commuting.*

households can attain their reservation utility level at lower levels of consumption. That is, people accept less consumption to get better weather. Nothing else about the model changes.

How does this affect the equilibrium city? As A increases and amenities get better, then; (1) equilibrium consumption falls, (2) the rent gradient intercept increases so rent goes up everywhere, and (3) the city grows in extent and population. This is illustrated in figure 1.3.

Sunny cities should be bigger and have higher rent than snowy ones, and the people in sunny cities should also consume a little less than people in snowy cities. The people in sunny cities are achieving the reservation utility level partly with sunshine and partly with consumption. The people in snowy cities must rely more heavily on consumption.

1.3.3 Comparative statics

We have a model that assumes: transportation is costly, everyone wants to work in the center, people arrange themselves so that no one wants to move, i.e., spatial equilibrium. These assumptions imply the downward sloping rent gradient that characterizes land markets almost everywhere, and that figure 1.1 illustrates for Japan and France.

Box 1.3.2: Partial differentiation

Given a univariate function $f : R \to R$, or $f(x) \in R$, we have

$$\frac{df}{dx} = \lim_{\epsilon \to 0} \frac{f(x+\epsilon) - f(x)}{\epsilon}$$

This is the "instantaneous slope" of f at x.

Partial differentiation is the generalization of this idea to surfaces. Consider a function $F : R^2 \to R$, or $F(x_1, x_2) \in R$. This function describes a surface, a height for each point in the plane. How do we think about the slope of such a surface? What we want is a tangent plane rather than a tangent line.

With partial differentiation, we think about the slope of such a plane along one axis. Thus, given $F(x_1, x_2)$, we define

$$\frac{\partial F}{\partial x_1} = \lim_{\epsilon \to 0} \frac{F(x_1+\epsilon, x_2) - F(x_1, x_2)}{\epsilon}$$

This is exactly analogous to the univariate derivative, if we imagine that we are finding the slope of a "slice" of the surface parallel to the x_1 axis.

Mechanically, treat the "other variables" as constant and use all the rules you know from univariate differentiation. For example, if $F(x, y) = 2x + 3y^2 + 2xy$ then $\frac{\partial F}{\partial x} = 2 + 2y$ and $\frac{\partial F}{\partial y} = 6y + 2x$. This should be in your calculus book.

It would also be nice to work out whether the model makes other predictions we can check. If the model makes predictions that are obviously not in line with reality, we will know we have a problem. With that in mind, we now consider some "comparative statics". That is, we ask how the monocentric city changes as we change, for example, wages or commuting costs while holding everything else fixed (we already looked at what happens when amenities change). Once this is done, we will have a collection of predictions for the model that we can compare to what we observe in the real world.

Changes in an open city as commuting cost, t, changes

If commuting costs fall, then utility and consumption stay the same in an open city. The household at $x = 0$ has a free commute, so its commute cost is unaffected by

Figure 1.6: Monocentric city comparative statics as commuting costs change

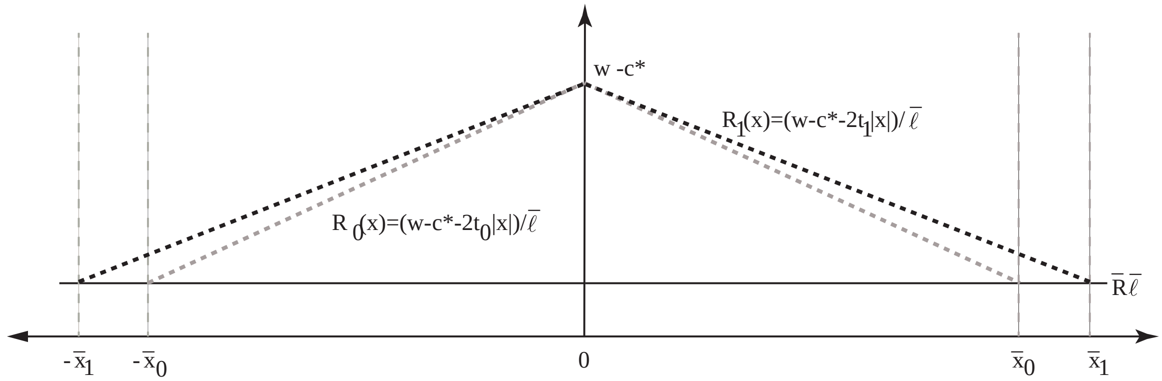

Note: *The dashed gray line describes an equilibrium land rent gradient in an open city when commuting costs are high, and the dashed black line describes an equilibrium land rent gradient in the same city when commuting costs fall. That is,* $t_1 < t_0$*. As commuting costs fall, land rent increases everywhere except at the* CBD *where the household's commute has length zero. As commuting becomes less expensive, at each location, the household has more income to divide between consumption and land rent. But the level of consumption is fixed by the reservation utility level. This means that the only thing that can adjust is the price of land. As land rent increases, it means that households can afford to bid a few marginal parcels away from farmers, and the edge of the city moves out from the* CBD.

the change in t, but households a little further from the CBD have the cost of their commute fall. Because the sum of land rent and commute cost is the same at all locations, this implies that the land rent gradient flattens, but its intercept stays the same. The only way consumption can remain constant at the reservation level, c^*, as commute cost falls is if land rent increases. At the edge of the city, land rent goes up and a few more households bid land away from farmers and the extent of the city increases. Because land rent increases everywhere, the total land rent paid to absentee landlords increases. Later on, we'll consider empirical evidence about this comparative static. For that purpose, it is helpful to note that as unit commute costs fall, a larger share of people live outside any fixed radius. This is all illustrated in figure 1.6

We can derive all of this analytically using partial derivatives (see box 1.3.2 if you need help with this). To proceed, take the partial derivative of the land rent gradient with respect to the commute cost,

$$\begin{aligned}\frac{\partial R(x)}{\partial t} &= \frac{\partial}{\partial t}\frac{w - c^* - 2t|x|}{\bar{\ell}} \\ &= \frac{-2|x|}{\bar{\ell}}\end{aligned}$$

This derivative is negative, so as t increases, rent falls at each x, and conversely (we're ignoring the corner where $R = \overline{R}$.) We can do exactly the same thing to see what happens to the extent of the city as commute costs change,

$$\begin{aligned}\frac{\partial \overline{x}}{\partial t} &= \frac{\partial}{\partial t}\frac{w - c^* - \overline{R\ell}}{2t} \\ &= -\frac{w - c^* - \overline{R\ell}}{2t^2} < 0.\end{aligned}$$

This derivative is negative, too, so as t increases, the length of the city falls.

Finally, we can check what happens to city population as commute costs change. Because $N = 2\overline{x}/\bar{\ell}$, we can use the chain rule to get,

$$\begin{aligned}\frac{\partial N}{\partial t} &= \frac{\partial}{\partial t}\frac{2\overline{x}}{\bar{\ell}} \\ &= \frac{2}{\bar{\ell}}\frac{\partial \overline{x}}{\partial t}\end{aligned}$$

and so the population of the city changes just like its extent. Thus, we obtain the same results analytically that we see in figure 1.6.

Changes in an open city as the wage changes

Figure 1.7 shows changes as wages increase in an open monocentric city. As wages rise, utility and consumption stay the same. This follows immediately from the open city assumption and spatial equilibrium. The slope of the land rent gradient is determined by the unit commute cost, and this also remains fixed. The intercept of the land rent gradient increases by exactly the amount of the wage increase, $w_1 - w_0$. This is the only way that we can balance the budget for households at $x = 0$ and keep consumption constant. The same increase has to occur everywhere, and for the same reason. Thus, an increase in the wage gives us a parallel shift up in the land rent gradient that offsets the wage increase. As a result, the extent of the city increases a little bit and and population increases. Aggregate land rent increases by almost the exact amount as the total wage bill.

It bears repeating that almost all of the benefit of an increase in wages is collected by absentee landlords. Even though wages go up, residents' consumption is

Figure 1.7: Monocentric city comparative statics as the wage changes

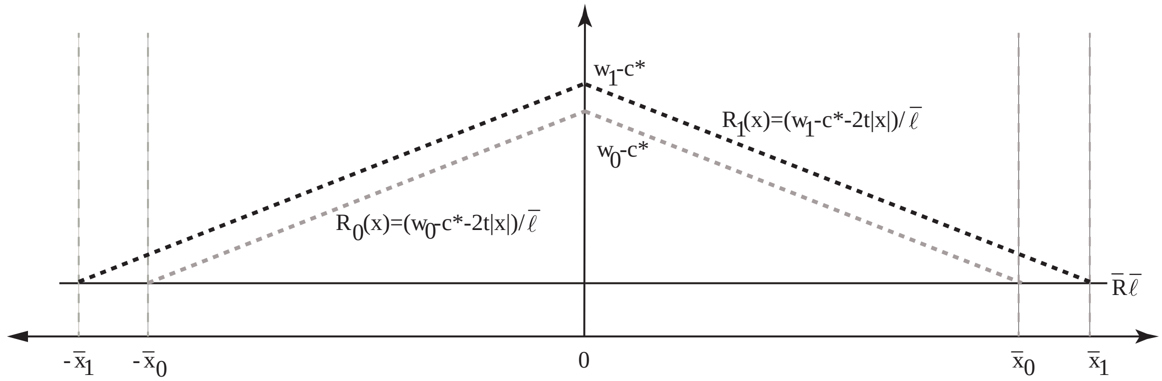

Note: *The dashed gray line shows an equilibrium land rent gradient for a low wage, and the dashed black line for a higher wage. As the wage increases in an open city, households everywhere see an equal increase in their income. Because consumption must stay constant, and because commute costs don't change, the only way to make sure the household budget balances if is land rent goes up by an amount that exactly offsets the wage increase. Because the value of land in residential use goes up everywhere, the extent of the city increases a little bit as a few more households are able to bid land away from farmers at the edge of the city.*

unchanged, so household budgets at any given location x can only balance if the wage increase is passed directly on to the landlords. At the time of this writing, there is a lot of policy interest in the US about increases to the local minimum wage. What does this comparative static suggest about the likely winner from these policies?

Changes in a closed city as commuting costs t change

Comparative statics in closed cities are quite different from those in open cities. Figure 1.8 illustrates what happens in a closed city when the cost of commuting increases.

Three things must stay fixed in a closed city as the cost of commuting increases. First, the size of the city cannot change. It is fixed by the fact that the number of people and per capita land consumption are both fixed. Second, land rent at the edge of the city cannot change because the rent required to bid land away from farmers also doesn't change. Finally, in a spatial equilibrium, the sum of land rent and commuting must be the same at all occupied locations, otherwise consumption differs across places and we don't have an equilibrium.

How can we satisfy these three conditions as the unit cost of commuting rises? If commute costs increase, the rent gradient must get steeper to keep the sum of rent and commuting constant. In addition, land rent must stay constant at $\bar{x}$. Together,

Figure 1.8: Monocentric city comparative statics as commuting costs changes in a closed city.

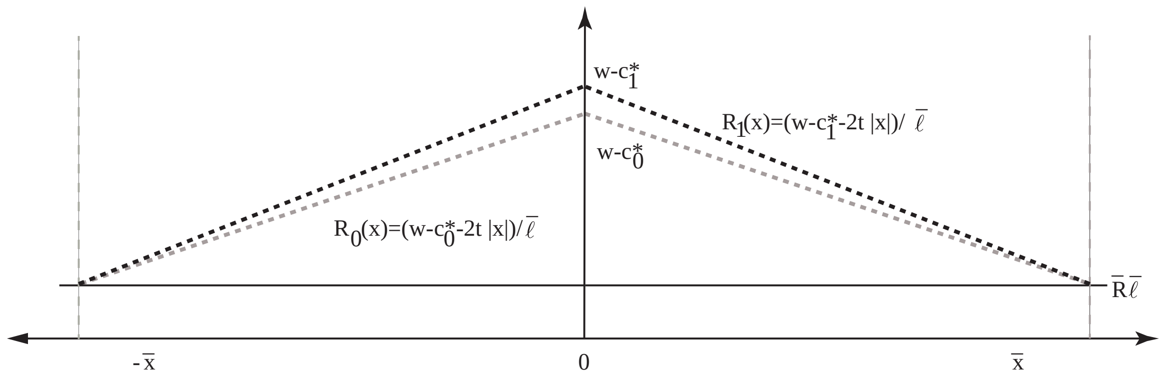

Note: *The dashed gray line shows the equilibrium land rent gradient in a closed city with low commute costs, and the dashed black line shows the land rent gradient in the same city when commute costs increase. In a closed city, the population is fixed, this fixes the extent of the city. With the extent of the city fixed, the land rent gradient must adjust so that the most remote household can just bid land way from farmers. As commute costs rise, land rent must increase in order to keep the sum of land rent and commute costs constant. In the closed city, it is the rent at the most remote location that is fixed by our assumptions. In contrast, for an open city, the reservation utility level fixed the level of land rent at* $x = 0$.

this means that rent at $x < \overline{x}$ must increase, and in particular, that the rent at $x = 0$ goes up. Because income is fixed, the increase in rent and commute costs means that consumption and utility falls. This is quite different from the open city where the utility of households is fixed, and changes end up affecting land rent without affecting household utility.

In an open city, the supply of people is perfectly elastic. All changes fall on landlords, good or bad. With a closed city, the supply of people is perfectly inelastic, so some of the change in commuting cost falls on the households. This highlights the importance of knowing how responsive is migration to local economic conditions for understanding the distributive consequences of urban policy. If we want to change something in a closed city, it affects the welfare of residents. In an open city, the welfare of residents is fixed, and payments to landlords change.

1.3.4 Land rent and welfare

In an open city equilibrium, each household gets $u(c^*) = \overline{u}$, and they get this payoff no matter how much rent they pay. In this sense, land rent is a measure of the benefit to

Figure 1.9: Aggregate land rent in the monocentric city

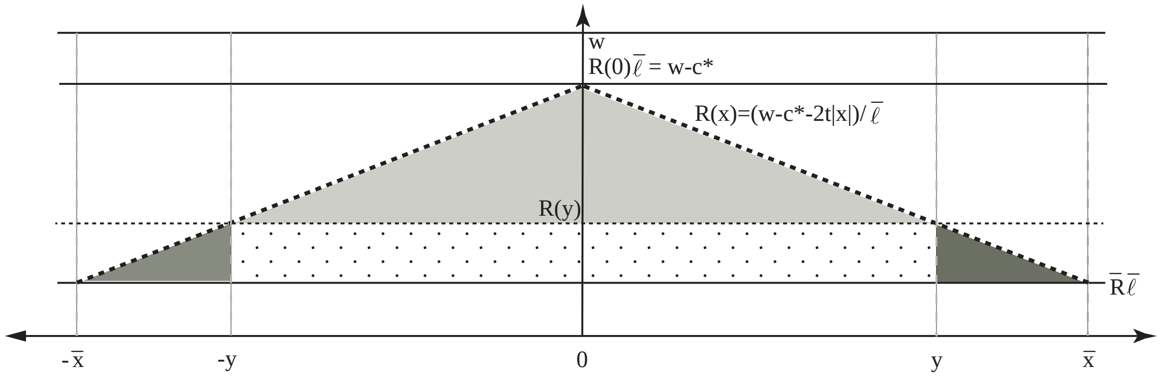

Note: *The dashed black line describes a land rent gradient for an open city.* $\overline{x}$ *is the edge of the city in equilibrium. A planner would like to choose an extent of the city to maximize aggregate land rent, taking as given that the city is open and households must satisfy the free mobility condition. If the planner chooses an extent of the city smaller than equilibrium,* $y < \overline{x}$*, then aggregate rent is less than for a city with edges at* $\overline{x}$*. In particular, the rent described by the two dark gray triangles is lost. If the planner chooses an extent greater than* $\overline{x}$*, then the planner must subsidize the marginal households to allow them to bid land away from farmers and still consume* c^**.*

a household from living in the city. They can get payoff $\overline{u}$ in the reservation location. In the city, they get this payoff and manage to pay land rent in addition. This suggests that aggregate land rent, the sum of land rent paid by all urban residents, is a measure of welfare. It is the collective willingness to pay to live in the city. It follows immediately, that changes in land rent indicate changes in welfare.

This is an important conclusion. Land rent is relatively easy to observe, much easier than utility levels, and so the fact that land rent measures welfare gives us a way to use easily observable data to think about the welfare implications of changes in the urban environment. Drawing conclusions about welfare from things that are easy to observe is not something that economists get to do very often. Indeed, in the next section we will show how we can use this intuition to value school quality or other place based attributes using data on real estate prices. There is an important caveat to this conclusion; it starts to break down once we start to think about models where not all households are the same. We'll return to this problem in Chapter 6.

Now that we have a way of measuring welfare in a city, it is natural to ask whether the equilibrium city maximizes welfare. A little more precisely, we have described how the monocentric city arises as an equilibrium outcome when everyone pursues their own narrow self interest. What would happen if a rent maximizing planner organized the city, subject to free mobility for the households? Would the resulting rent maximizing city be different from an equilibrium city?[3]

In our optimal city, we still allow free mobility, so, as in the equilibrium city, we must have

$$w - c^* = R^*(x)\bar{\ell} + 2t|x|.$$

Rearranging, we have

$$R^*(x) = \frac{w - c^* - 2t|x|}{\bar{\ell}}$$

at all occupied locations. Given this, our planner wants to choose the extent of the city, y to maximize total land rent, taking as given that rent is given by this expression at any location in urban use.

To avoid an involved calculus problem, figure 1.9 makes the argument graphically. This figure is like those we have used throughout the Chapter to illustrate the monocentric city model. In particular, $\bar{x}$ is the equilibrium extent of the city. We would like to consider whether a planner choosing the extent of the city to maximize aggregate rent would choose something different.

When the planner chooses $\bar{x}$ as the extent of the city, then aggregate rent is just what we would have for the equilibrium city. It is the sum of the light gray, dark gray and dotted regions under the land rent gradient. Suppose our planner chooses a slightly smaller extent for the city, $y < \bar{x}$? Then aggregate land rent is just the area under the land rent gradient between $-y$ and y. This is the sum of the light gray and dotted regions. This is clearly less than the aggregate land rent that results if the planner chooses $\bar{x}$ as the extent of the city. Now what if the planner chooses $y > \bar{x}$. Then the marginal increase in urban land rent does not offset foregone agricultural land rent. In fact, the planner has to subsidize the marginal urban resident in order to allow them to bid land away from a farmer and still afford enough consumption that they don't want to move away. It follows that the monocentric city that emerges in equilibrium is "optimal" in the sense that it maximizes land rent.[4]

1.4 Application #1: Learning the value of school quality from real estate prices

An interesting implication of the monocentric city model is that land rent can never be discontinuous. To see this, imagine the rent gradient drops discontinuously as we move

[3]You might recognize the parallel between this question and the one answered by the first fundamental theorem of welfare economics. This theorem states that, under weak conditions, if a market equilibrium exists then it is Pareto optimal. We will find something similar here.

[4]This is slightly weaker than the first welfare theorem because rent maximization is implied by Pareto optimality, but not conversely.

Figure 1.10: School district boundaries in Melrose Massachusetts around 1990

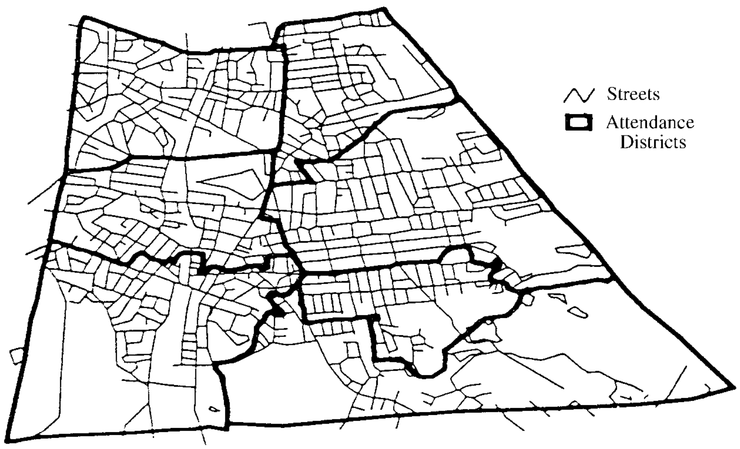

Note: *Heavy black lines show school attendance zone boundaries in Melrose Massachusetts around 1990. Lighter lines are streets. Reproduced from Black [1999].*

away from the CBD. In this case, the household at the high side of the discontinuity can move to the low side, experience almost zero change in commute costs, and a discrete drop in rent. This contradicts the idea that this was an equilibrium to start with. A household can move and make themselves better off. A similar argument works if there is a discontinuous increase in land rent.

This means that we can have a discontinuous rent gradient only if amenities vary discontinuously. In this case, spatial equilibrium requires that rent vary discontinuously in order to equalize utility across locations. This intuition motivates the "border discontinuity design" for learning about the value of amenities that vary discretely as we move across the landscape.

Black [1999] uses this idea to examine the value of school quality by looking at how housing prices vary when we cross a school district boundary where school quality varies. She considers the relationship between school quality and real estate prices between 1993 and 1995 for three counties in Massachusetts. Figure 1.10 illustrates this geography for a single city.

Black matches data describing elementary school average test scores (a proxy for school quality) and real estate transaction data to this map. School quality varies discretely at an attendance zone boundary. How much is this worth? The logic of spatial equilibrium tells us that as long as nothing else changes at an attendance zone boundary, the land price gradient should be continuous. If we see a jump, it must mean that the value of the properties is changing to reflect the different value of attending schools in the different attendance zone.

To measure this gap in real estate prices (if it is present), Black restricts attention to transactions within a few hundred yards of a school attendance zone boundary, and estimates the following regression,

$$\ln(\text{House price}_i) = A_0 + A_1 \text{test score}_i + A_2 \text{border indicators} + \text{controls}_i + \varepsilon_i$$

The parameter A_1 tells us the size (in log points) of the change in house prices at the border. If real estate markets are in "spatial equilibrium" this should tell us the value of improving test scores. Black finds that A_1 ranges between 0.013 and 0.031, so a 1 point increase in test scores increases the logarithm of housing prices by between 0.013 and 0.031, which works out to between a 1-3% increase in housing prices. In Black's sample, about 90% of all houses lie in attendance zones with test scores between 25.2 and 29.8, so moving from the 10th to the 90th percentile of school district quality results in an increase in a house price increase of between about 4 and 12%.

It's worth taking a minute to think about how neat this is. Suppose you did not know about this trick, how would you go about figuring out the value of an improvement in test scores? It would likely involve tracking what happened to students who were otherwise similar, but went to better and worse schools, trying to figure out how their lives turned out, and then trying to attach a dollar value to this difference. This border discontinuity design using real estate prices is much simpler. It lets us work out the value of better schools in one step.

It's also worth noting the problems with this method. First, we're getting the value of school quality to the parents not to the students. If you think parents don't value their childrens" education the way they should, then this could be a problem. Second, it is possible that school attendance zones follow features of the landscape that divide the nice places from the unpleasant. For example, they might follow rivers where one bank is swampy and the other is not. This is a well known problem with these sorts of border discontinuity research designs, and Black follows good practice and carefully excludes attendance zone boundaries where this sort of problem might obviously arise. Third, it may be that what this exercise is picking up is not the value of better schools at all, but the value of living near people who value better schools. In a similar exercise done a few years after Black's study, Bayer et al. [2007] found that people living on the high score side of a school district boundary had higher incomes and were more likely to be college educated and white. Like Black, Bayer et al. find that real estate prices are higher on the high score side of a boundary, but the fact that people are sorting themselves into better and worse school districts on the basis of other characteristics means that we can't rule out the possibility that part of what people value is proximity to affluent, college educated white people, not access to better schools. We consider such sorting in more detail in Chapter 10.

Figure 1.11: Plot of relative tax rate versus relative house price in for the US 2000-2016

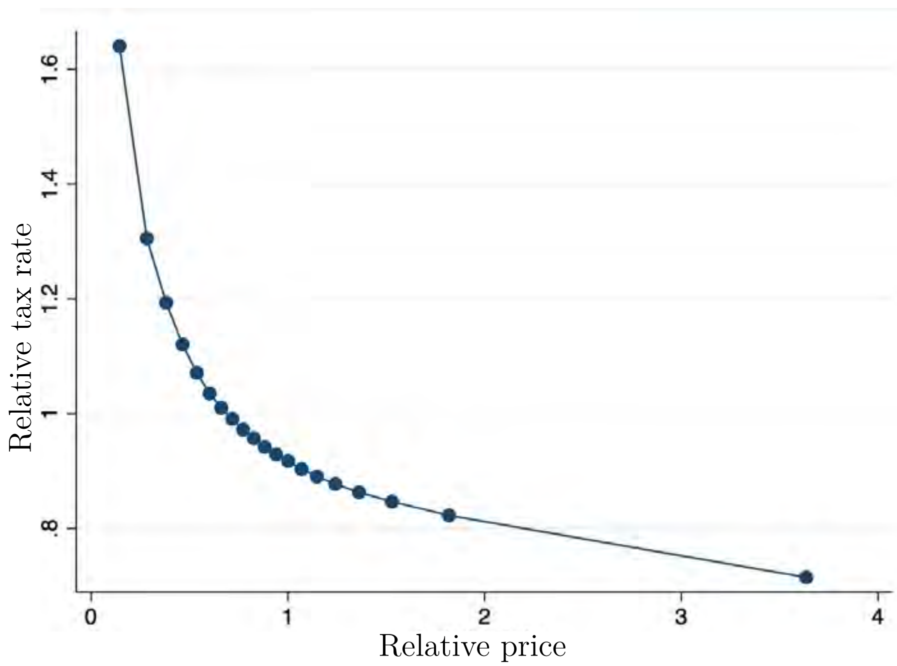

Note: *Relative price is property sale price divided by jurisdiction average price in year of sale. Relative tax rate is property's tax rate divided by jurisdiction average tax rate in the year of sale. The tax rate is the tax due in the year of sale divided by the sale price. Binned scatter plot shows average relative tax rate and average relative price by 20 quantiles of relative sale price. Based on 26 million residential sales. Figure and figure note reproduced from Berry [2021].*

1.5 Application #2: Detecting racist property tax assessments

Consider the problem of unfair property tax assessments in Chicago. The June 20, 2024 edition of the *Chicago Tribune* reports that

> `An unprecedented analysis by the Tribune reveals that for years the county's property tax system created an unequal burden on residents, handing huge financial breaks to homeowners who are well-off while punishing those who have the least, particularly people living in minority communities.`
>
> `The problem lies with the fundamentally flawed way the county assessor's office values property.`

> `The valuations are a crucial factor when it comes to calculating property tax bills, a burden that for many determines whether they can afford to stay in their homes. Done well, these estimates should be fair, transparent and stand up to scrutiny.`
>
> `But that's not how it works in Cook County, where Assessor Joseph Berrios has resisted reforms and ignored industry standards while his office churned out inaccurate values. The result is a staggering pattern of inequality.`[5]

The figure 1.11 illustrates the extent of the problem using a national sample of real estate transactions and property tax bills. This figure is based on data describing the sales price on the x-axis, and the ratio of the property tax bill to the sale price on the y-axis. Both quantities are adjusted statistically for differences in averages across counties and school districts. The steep downward slope means that less expensive houses are paying a greater share of house value as taxes. The people who live in less expensive houses, people who are disproportionately black and Hispanic, have a bigger property tax bill relative to the price of their houses than the disproportionately white people who live in more expensive houses.

This looks bad for tax assessors nationwide. However, while it is plausible (and even likely) that there has been misbehavior by tax assessors, at least in Chicago, this figure is not the smoking gun it first appears.

In particular, if property taxes are based on market prices and market prices capitalize tax assessments, the relationship we see in figure 1.11 is just what we would expect when assessors are behaving fairly. The argument has three steps. First is looking at how property rental prices and property taxes are related. Second is looking at how rental and asset prices are related, and third is putting the first two together.

1.5.1 Property taxes and rental prices

Consider a monocentric city and suppose that land is subject to a property tax rate τ. How does this change the equilibrium? For the purpose of this problem, it's going to be important to discriminate between economic rent and contract rent. Recall economic rent is the whole value of the property to the tenant, and contract rent is what the tenant pays to the landlord. Let R_C be the contract rent in the taxed city, and suppose that the household at x pays property tax $\tau R_C(x)$, so the household's total payment for the property is $(1+\tau)R_C(x)$ every month. Then the household's

[5] `https://apps.chicagotribune.com/news/watchdog/cook-county-property-tax-divide/assessments.html`, September 20, 2024.

Figure 1.12: Contract rent and economic rent in an open city with a property tax

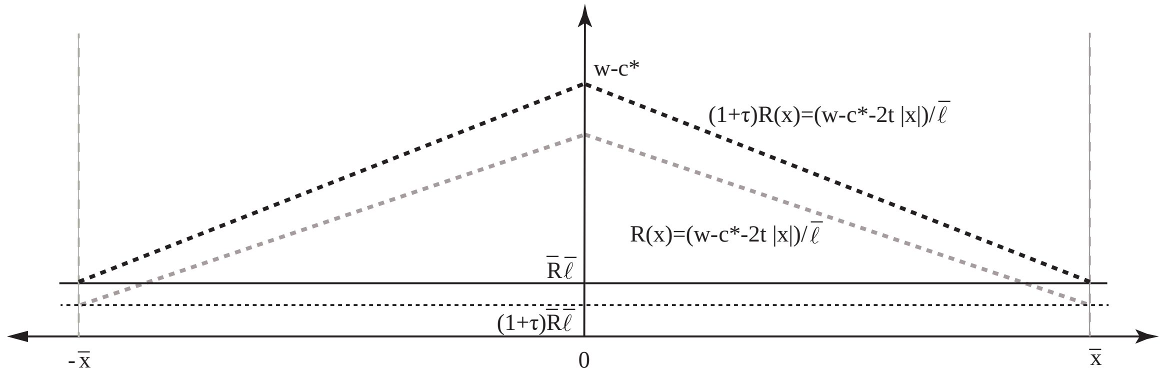

Note: *The dashed black line gives the land rent in a monocentric city without a property tax. This is the "economic rent gradient" and it is exactly the same as we have seen in other open city examples. The dashed gray line is the "contract rent gradient", what the tenant pays the landlord, before paying a property tax. The difference between the blue and green lines is the tenant's tax payment.*

problem is

$$\max_{c,\, x} u(c) \tag{1.16}$$
$$\text{s.t. } w = c + (1+\tau)R_C\bar{\ell} + 2t|x|.$$

This is still an open city, so we should have $u(c^*) = \bar{u}$ at all occupied locations. This requires constant consumption of $c^* = u^{-1}(\bar{u})$, just as in a city without property taxes.

Substituting c^* into the budget constraint and rearranging, we get

$$(1+\tau)R_C = (w - c^* - 2t|x|)/\bar{\ell}.$$

If we compare this expression to the expression for the land rent gradient in an untaxed city in equation (1.15), we see that the sum of the contract rent and taxes is exactly equal to the economic rent in a city without a property tax. Contract rent plus taxes sums to economic rent. In math, we have

$$R^*(x) = (1+\tau)R_C(x). \tag{1.17}$$

This means that adding taxes to the household's problem does not change anything about the city, except that some of the money that would have been collected by absentee landlords is collected by the government, exactly the same conclusion we reached in section 1.1.

Figure 1.12 illustrates. The dashed black line in this figure gives the land rent gradient in the untaxed city. This is the "economic rent gradient". After households pay this amount of rent and pay for their their commute, they are just able to purchase the reservation consumption bundle, c^*. The dashed gray line describes the "contract rent gradient". This is what the household pays the landlord. This is just enough below the dashed black line, that the tax payment makes up the difference. Because the household doesn't care whether it pays the city or the landlord, think back to our example of the friendly gangster in section 1.1, the household makes decisions on the basis of contract rent plus taxes. That is, on the basis of economic rent. But this means that nothing about the household's decision changes.

This is a pretty surprising result. It says that property taxes don't change behavior at all. That bears repeating. Property taxes don't change behavior at all. This is a special feature of property taxes.

To understand, why this is important, consider the problem of a legislature that needs to raise 100$ per person in tax revenue. It has the choice of a tax which simply collects 100$ from everyone, or a tax which collects 100$ from everyone who stands on one foot for a minute on April 15, and 200$ from everyone else. We expect both systems of taxation to raise the same revenue, but one makes the average taxpayer worse off by whatever discomfort they endure standing on one foot for a minute. That is, the second tax system creates an incentive for tax avoidance behavior, and tax avoidance behavior is usually wasteful. People engage in it not because they like it, but because it reduces their tax burden.

Taxes that raise revenue without creating an incentive for avoidance behavior are rare, and they are special because they allow the government to raise one dollar of revenue at the cost of only one dollar of harm to the taxpayer. With avoidance behavior, the cost of a dollar of revenue is always more. It is one dollar plus the cost of the avoidance behavior. A city with a property tax is full of people who act in exactly the same way as they would in a city without a property tax. In this sense, a property tax is at least as good as any other way of raising government revenue.

This result is widely known as the "Henry George Theorem" (although it is not easy to find it stated explicitly anywhere, in particular in Henry George's writings.) Two caveats apply. First, if people are not all identical, this result starts to break down. Second, it's important to also tax agricultural land. Otherwise there is an effect on the extensive margin. Some people move away from the edges of the city because untaxed farmers outbid them for land.

In theory, the way property taxes are assessed is as follows. First, an assessor assigns your house a value. Usually, this value is deliberately close to the house's market value.[6] Second, the municipal government chooses a "mill rate", typically around 1%.

[6] If you ever own a house, you will invariably conclude that the assessor thinks your house is much

Each homeowner's tax bill is the product of the mill rate and their assessed value. In practice, there are lots of opportunities for unfairness and malfeasance, but for the present purpose, we'll suppose that this has a small impact on tax rates.

We've finished working out how property taxes and property rental prices are related. However, property prices depend on *asset prices* not *rental prices*, so knowing the relationship between property taxes and rental prices is not enough to understand the whole process. Our next step is to figure our how rental prices are related to asset prices.

There is a caveat to this argument. If we are being precise, we have so far considered the sale of land, rather than the sale of houses. Property taxes are almost always collected as taxes on the joint value of land and any structure on the land. This means that "property taxes" are also a tax on houses. This matters because it creates an opportunity for avoidance behavior. "Over-taxed" houses transact for less money, and their owners write larger checks to the city each year. The problem is that you pay property taxes on improvements to your house, too. If you add a room, you pay property tax on the value of this addition forever. If you are subject to a higher property tax, home improvements cost more. This means that a high property tax disincentivizes home improvement and maintenance, or said another way, incentivizes blight.

1.5.2 Land rent and capitalization

How are rent and asset prices related? To answer this question, we need to work out the mathematics of "discounted present values". Let ρ be the real interest rate. One dollar today turns into $1 + \rho$ in a year. P is the purchase price of a property and R the rental price for one year. If $\rho P < R$ then renters should buy their properties and pocket the difference. If $\rho P > R$ then owners should sell and become renters. Only when $\rho P = R$ is there no opportunity for intertemporal arbitrage. So, we should have $\rho P = R$. That is, rent equals one year of interest on the asset price of the property. Box 1.5.1 gives another derivation of this same result.

1.5.3 Fair assessment of property taxes

What does all of this mean for the relationship between property taxes and the sale price of houses?

---

more valuable than anyone else in the world. On the other hand, California's infamous Proposition 13 prevented changes in assessed value except when a property changes hands, this means that the assessed value of properties that have not sold for a long time are often much lower than similar houses that have changed hands more recently.

Restating equation (1.17), we have,

$$\begin{aligned} R^*(x) &= (1+\tau)R_C(x) \\ &= \tau R_C + R_C(x). \end{aligned} \tag{1.21}$$

Now we need some notation. Let $V(x)$ be the "economic asset price". That is, the discounted present value of economic rent $R^*(x)$, and let $V_C(X)$ be the "contract asset price", that is, the discounted present value of contract rent.

Starting from equation (1.21), we have

$$\sum_{t=1}^{\infty} \delta^t R^*(x) = \sum_{t=1}^{\infty} \tau \delta^t R_C + \sum_{t=1}^{\infty} \delta^t R_C(x),$$

or

$$V(x) = (1+\tau)V_C(x).$$

If we take logs of both sides, and recall that $\ln(1+x) \approx x$ for x small, we get

$$\begin{aligned} \ln V(x) &= \ln(1+\tau) + \ln V_C(x). \\ &\approx \tau + \ln V_C(x). \end{aligned}$$

Rearranging, we get

$$\tau \approx \ln V(x) - \ln V_C(x). \tag{1.22}$$

This is the punchline. Notice that in a city with a property tax, we will never observe $V(x)$. These are the transaction prices that would occur in the (counterfactual) absence of a property tax, but we will observe the tax rate and $V_C(x)$.

Now suppose we conduct a regression of the tax rate on observed asset prices. What would this look like? Letting i index transactions, it will be something like this,

$$\tau_i = A_0 + A_1 \ln V_C(x)_i + \varepsilon_i. \tag{1.23}$$

From equation (1.22) expect that A_1 would be about -1. If we were going to plot this, it would show a rapid decrease in the tax rate with value of the property, exactly what we see in figure 1.11. The current system of property tax assessment in Chicago, or in the US as a whole may well be terribly corrupt and unfair, but figure 1.11 does not make this case. That the tax rate declines with property price is an implication of the way that property taxes are capitalized into property prices.

This is a dramatic and, to me, surprising result. Why does it work? We know that property prices affect property taxes. This is given in the rules for how property taxes are calculated. But property taxes also affect property prices. This is the logic of capitalization. This means that the relationship we see in figure 1.11 has to reflect both of these relationships. The math we've just worked out shows how these two relationships work together to create a downward sloping relationship between the tax rate and transaction price, with no racism required.

1.6 Conclusion

We've now developed the basic version of the monocentric city model pretty thoroughly.

This model assumes: spatial equilibrium, costly commuting, and central employment. The open city model makes the following predictions.

First, $R^*(x)$ decreases in x. We've seen that this is correct, and there is more evidence on this point to come.

Second, as commuting costs, t, decrease, utility, $\overline{u}$, stays constant by assumption. Constant utility immediately implies that equilibrium consumption, c^*, also stays constant. This, in turn requires that; the rent gradient gets flatter and its intercept stays the same, and that the extent and population of the city increases. We will see some evidence about this later.

Third, as wages, w, increase, utility and consumption, $\overline{u}$, c^* stay constant (by assumption). The slope of the rent gradient is unchanged, but its intercept increases by the same amount as the wage increase. The extent of the city and its population both increase, and aggregate rent increases by about the same amount as the aggregate wage bill.

We haven't worked out what happens as agricultural rent changes. This is straightforward, but there is not much empirical evidence about this comparative static, so I am going to leave this as an exercise.

As amenities, A, increase, utility stays constant, but consumption c^{**} falls. The slope of rent gradient is unchanged, and the intercept increases. The city gets longer, population and aggregate land rent both increase.

Changes in property taxes do not change anything except how much rent is collected by absentee landlords. This is called the Henry George Theorem.

Spatial equilibrium requires that rent gradients be continuous, unless something that people value about the location changes discontinuously. This intuition gives rise to the widely used border discontinuity research design for evaluating location specific attributes.

This is a good start, but leaves open a few questions. First, the shape of the rent gradient that the model predicts is wrong. The model is predicting a linear rent gradient when in reality it decreases much faster than this. In its current form, the monocentric city model is a model of land allocation. Adding a description of housing (as opposed to just land) in Chapter 3 will help with this. Second, why are people in the center? This is a central assumption. Implicitly, there is a mill or big factory in the CBD where people are more productive than if they work elsewhere. So far, we've just assumed that people want to be in the center. It would be nice to understand a little bit more about why. We'll come back to this when we talk about agglomeration economies in Chapter 7.

Finally, the assumption that cities are all monocentric is contradicted every time we set foot in a suburban big-box store. Indeed, the assumption that people work only in the CBD is so obviously at variance with observation as to bring the usefulness of the monocentric city model into question. There are two responses to this. First, is to point to the body of empirical evidence confirming the predictions of the monocentric city model. The next chapter describes this evidence. Second is to work with models that allow both firms and households to choose their locations. This line of investigation is very technically demanding and was pioneered in a pair of papers, Fujita and Ogawa [1982] and Ogawa and Fujita [1980], and later refined and extended in Lucas and Rossi-Hansberg [2002]. More recently, it has been the subject of a literature on Quantitative Spatial Models, which are the subject of Chapter 6.

Box 1.5.1: Discount present value calculations

There is a second way to work out the relationship between rental and asset prices. It's a little more complicated, but it gives a better intuition about how capitalization works. Suppose the rent on a property is R every year, forever. The sales price is the value today of this stream of payments. R in one year is worth $R_1 = \frac{1}{1+\rho}R$ today. R in two years is worth $R_2 = \frac{1}{(1+\rho)^2}R$ today, and so on. R every year forever, starting in one year is worth

$$\begin{aligned} V &= \frac{1}{(1+\rho)}R + \frac{1}{(1+\rho)^2}R + \frac{1}{(1+\rho)^3}R + ... \\ &= \sum_{t=1}^{\infty} \frac{1}{(1+\rho)^t}R. \end{aligned} \tag{1.18}$$

To evaluate this, start by defining $\delta = \frac{1}{(1+\rho)}$ (δ is called the "discount factor"). We can now rewrite equation (1.18) more compactly as,

$$V = \sum_{t=1}^{\infty} \delta^t R. \tag{1.19}$$

Multiplying both sides by δ we get,

$$\delta V = \delta \sum_{t=1}^{\infty} \delta^t R. \tag{1.20}$$

Subtracting equation (1.20) from (1.19),

$$\begin{aligned} V - \delta V &= \sum_{t=1}^{\infty} \delta^t R - \delta \sum_{t=1}^{\infty} \delta^t R \\ \implies (1-\delta)V &= \delta R + \delta^2 R + \delta^3 R + ... \\ &\quad - \delta^2 R - \delta^3 R - \delta^4 R - ... \\ &= \delta R \end{aligned}$$

Substituting in the definition of δ and rearranging, we get $\rho V = R$. That is, the rental price of land is equal to the interest payment on the asset price. The two ways of figuring out how rent and asset price are related are equivalent (this is pretty neat).

Problems

1. In this problem, we will work through an example of the monocentric city model. Assume we have a linear, open city. Let w=3, $\bar{l} = 1$, $p_c = 1$, $\overline{R} = 0.5$, $\overline{u} = 0$, and $A = 1$. Let $u(c) = \ln(c - 1)$.

   (a) Set up the household's problem. Assume we are in a spatial equilibrium, so everyone is optimizing and no one wants to move. Call consumption in this equilibrium c^*. What is $u(Ac^*)$ equal to?

   (b) Find c^*.

   (c) Using the constraint from the household's problem, find an expression for $\overline{x}$ in terms of $w, c^*, \overline{R}, \bar{l}$ and t.

   (d) Use the assumption that there is one unit of land at each x to derive an expression for N^* in terms of $\overline{x}$ and $\bar{l}$.

   (e) Use the household's equilibrium budget constraint and the equilibrium extent of the city to solve for the equilibrium rent gradient, $R^*(x)$.

   (f) Take derivatives of your expressions for $\overline{x}, N^*,$ and $R^*(x)$ with respect to t. How do the city extent, population, and equilibrium rent gradient change as transportation costs increase? Provide some intuition.

   (g) Assume that transportation costs increase from $t_0 = 1$ to $t_1 = 2$. What is the boundary of the city now? What is $R^*(0)$? Use these three points to draw a picture of how the rent gradient changes when t increases. Please label $R^*(0), \overline{R}$ and $\overline{x}$.

   (h) How would total land rent within the boundaries of the city change if we go from $t_0 = 1$ to $t_1 = 2$?

2. In this problem, we will analyze property taxes in the monocentric city model.

   (a) Assume we have an open, linear city with property tax rate τ_0. $R_0(x)$ is the land rent in this city. Set up the household's problem (you don't need to solve it).

   (b) Assume the tax rate increases from τ_0 to τ_1, where $1 + \tau_1 = (1.10)(1 + \tau_0)$. Set up the household's problem with this new tax rate.

   (c) Using what you know about c^* in an open city equilibrium, solve for $R_1(x)$ in terms of $R_0(x)$. How does the sum of rent and property taxes change?

   (d) Suppose landlords are responsible for paying the property tax. What does this suggest about the relationship between what tenants pay and property taxes?

3. In this problem, we will examine rental gradients in practice. Using Zillow or some similar real estate website, pick a radial road out from the center of Providence (for example, along Angel Street from Kennedy Plaza in Providence) and plot the prices of at least 15 similar properties as distance to the center increases. What do you find? You can do this for another city if you would like.

Chapter 2

The Monocentric City Model versus Data

The language around hypothesis testing in econometrics and statistics is both awkward and particular. We check if "the data fails to reject" a particular hypothesis. We're not trying to conclude that we're right, only that we're not clearly wrong. Gravity is just a theory. It explains the way apples fall from trees pretty well, but we're still keeping an open mind. The humility implicit in this awkward turn of phrase, "fail to reject", has always struck me as central to the scientific method. Our models, like the monocentric city model are always simplified, stylized descriptions of the world, and while they may make a lot of good predictions, eventually, we will find some sort of economic dark matter that forces us to revise or abandon any model.

With that said, the monocentric city model holds up better than just about any economic model I know. Even though it is highly stylized, cities don't really exist on a featureless line or plane and not everyone works in the center, it is hard to reject any of the comparative statics we worked through in Chapter 1 on the basis of the available data.

This Chapter surveys the evidence for this claim. In particular, the monocentric city model makes predictions about what should happen to land rent, and the extent and size of the city as transportation costs, amenities, and property taxes change. We here compare these predictions with the available empirical evidence.

2.1 Cities in real life

To check the predictions of the monocentric city, we need some real group of people to try to match up with the theoretical city. If you think carefully about this, it's pretty hard. The monocentric city model is an abstraction from reality. Finding some

Figure 2.1: MSA boundaries and lights at night in the Northeastern US

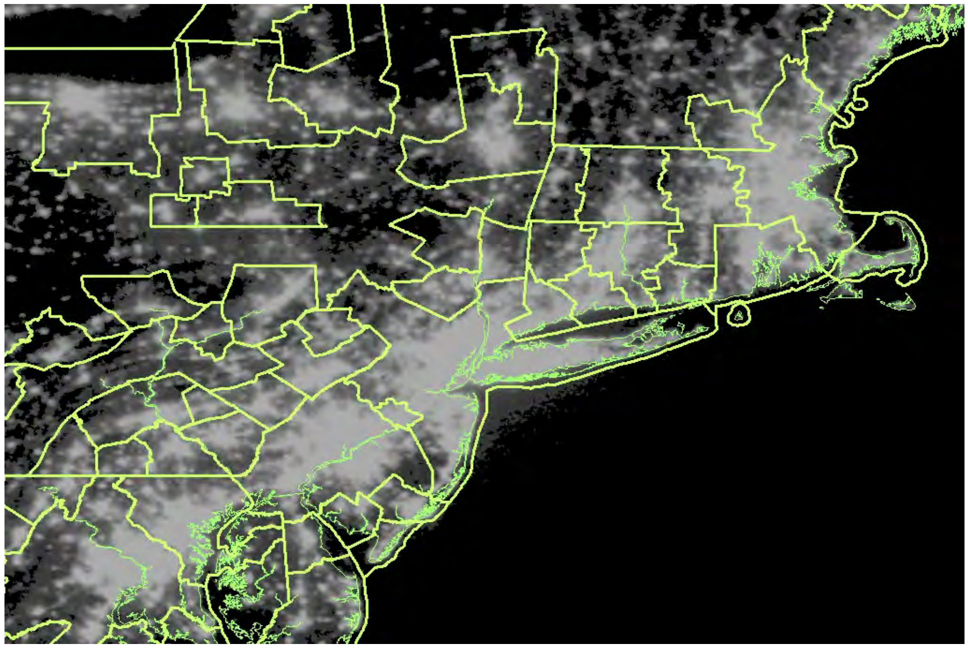

Note: *MSAs in New England in 2019 and lights at night in 2013. The New York MSA is in the center of the picture. If you look carefully, you will see that not everywhere is part of an MSA.*

real object that matches it closely is going to be hard and there is not going to be a single "right" answer.

A central feature of the monocentric city model is that it is a labor market. People work in the center if and only if they live in the city. This suggests that we want to be looking at some real world area that seems like a "labor market" in this same sense; an area in which most residents also work and where not many non-residents work.

In the US, the unit that satisfies these criteria most closely is probably the "metropolitan statistical area", or MSA. MSAs are metropolitan areas of at least 50k people, built from counties. They are purely reporting units and are defined by the US Census Bureau. There are a few different flavors, "micropolitan statistical areas", "core based statistical areas" (CBSAs), "consolidated metropolitan statistical area" (CMSAs). Definitions are easy to find on the census website. In every case, the idea is to build a metropolitan labor market from counties. Many of the empirical papers we discuss will use MSAs, or one of their variants, as the definition of "city". Figure 2.1 illustrates MSA boundaries in 2019 for the Northeastern US. The background shows lights at

Figure 2.2: Prefectural boundaries and lights at night in Northeastern China

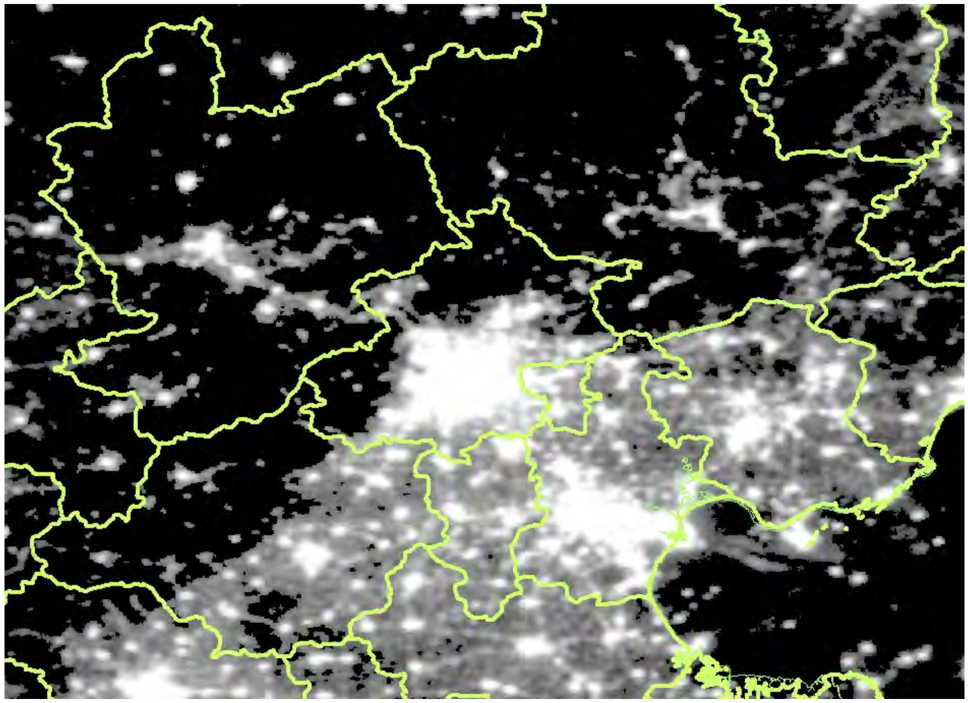

Note: *Prefectural cities in China in 2005. and lights at night in 2013, Beijing is central. Prefectural cities are the nearest analog to US MSAs. But, prefectures are also administrative units in China, while MSAs are purely reporting units in the US.*

night (from 2013) to show the extent of development. Notice that MSAs often contain a lot of undeveloped land. The dividing line between one MSA and the next is sometimes not matched with a dividing line in night lights, so while MSAs are probably the best empirical analog to the monocentric city that we have available, it is clear that reality is more complicated than our simple model allows.

Besides MSAs, there are many other candidate geographies that describe "cities". Municipal boundaries come immediately to mind. That is, the administrative boundaries of a city. This would be a natural unit to consider if we were thinking about issues related to municipal finance or services, but for thinking about the implications of the monocentric city model, they seem like a bad fit. Many municipalities do not contain the employment center where most of their residents work. Conversely, many of the workers in center city municipalities do not live in the municipality where they work. Many people who live in suburban Rye, NY, just north of Manhattan, work in Manhattan, not Rye. Conversely, many of the people who work in Manhattan do not live in New York City. The municipality is not the geography that the monocentric city model describes.

The census also describes "urban areas". While the census definition of urban area stretches over several paragraphs, the goal is to delimit the places where people reside, and the boundaries of urban areas often line up closely with remote sensing data showing built up areas. The boundaries of urban areas don't tell us anything about commuting patterns within the urban area. It follows that, they are also not obvious real world analogs to the geography of the monocentric city model.

Each of the geographies described above is based on US census definitions. Other countries often keep track of pretty similar units, either based on administrative or reporting boundaries. For example, in China, the unit that corresponds most closely to an MSA is the "Prefectural City". China covers about the same land area as the US and, like the US, has about 3000 counties and about 30 provinces, so US and Chinese counties are about the same size and Chinese provinces are a little larger than US states. However, while there is no administrative unit between county and State in the US, in China, the "prefecture" is an administrative unit in between county and province, with each prefecture being made up of a collection of counties in the same province. In general, counties are grouped in prefectures so that a prefectural government oversees a metropolitan area (though there are some rural prefectures), and these prefectures are the closest Chinese geographical unit to the US MSA. Figure 2.2 illustrates 2005 prefecture boundaries in Northeast China with nightlights for 2013 in the background. Beijing is in the center of the figure. Notice that, like the US, there is a lot of sparsely populated area in many of the prefectural cities.

Prefectural cities are unlike US MSAs in two regards, First, every Chinese county is part of a prefecture, while many rural counties in the US are not part of an MSA. Second, the prefecture is an administrative unit in China, but the US MSA does not have any administrative authority. It is a geography created solely for reporting purposes. Recalling that we want our empirical unit to describe a labor market, no one commutes outside the area boundary to work, it is not clear whether Chinese prefectural cities are better or worse analogs of the monocentric city model than US MSAs.

2.2 Rent gradients and Covid

Among its many effects on our lives, the Covid pandemic that began in 2020 reduced commuting costs for a large part of the workforce. In response to the pandemic, many people were able to work from home at least part of the week, which meant fewer commute trips. In the language of the model, this means that t decreased more-or-less in proportion to the decrease in commute trips. We saw in section 1.3.3 that a decrease in t implies that the rent gradient flattens, its intercept stays the same, and

Figure 2.3: Pandemic asset price and rent growth in New York City and San Francisco from December 2019 to December 2020

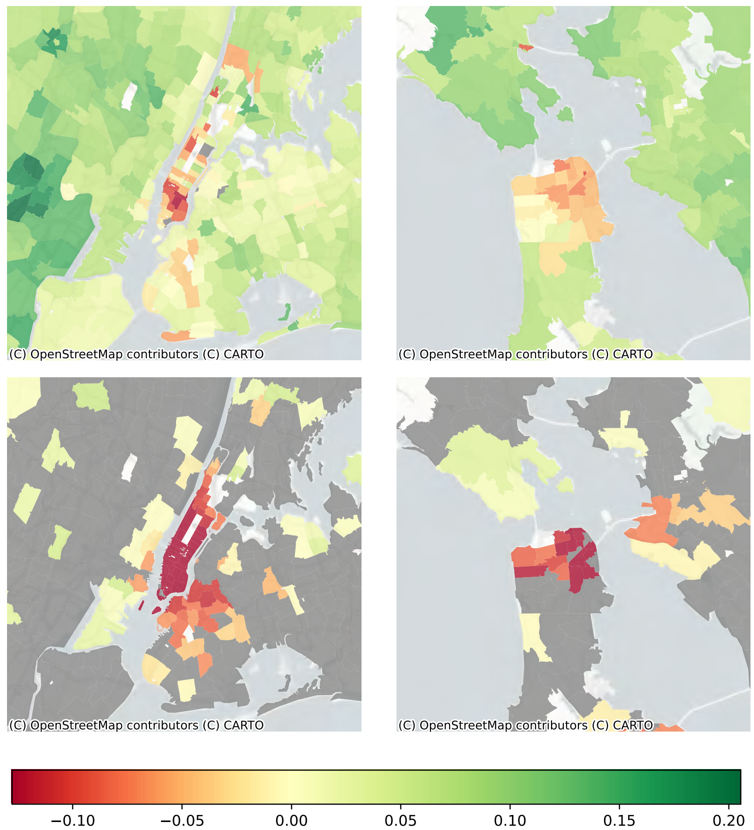

Note: *Year-over-year changes in prices (top two panels) and rents (bottom two panels) at the ZIP code level for the New York and San Francisco* MSAs *from Dec 2019 to Dec 2020. The bottom two rows zoom in on the city center. Darker green colors indicate larger increases, while darker red colors indicate larger decreases. From Gupta et al. [2021].*

the city gets longer and more populated.

The Covid pandemic also changed the amenity value of living close to the center of the city. Prior to the pandemic, a central residential location meant access to restaurants, bookstores and violin lessons. During the pandemic, the center of the city became a minefield of opportunities to contract an infection. Because the risk of infection was increasing with the number of people around, the risk of infection was lower in suburbs and rural areas, and so the second effect of the pandemic was to increase the amenity value of the suburbs relative to the center city. This is not quite the case we looked at (we had changes in amenities the same everywhere), but it is close. We expect this sort of change in amenities to have different on land rent in the center than the suburbs. It is going to decrease rent in the center. In the suburbs it will decrease rent by less than in the center, or even increase it. Adding the two effects of Covid, one operating through commute costs, and one through amenities, we should see urban land rent gradients flatten and decrease at the center. The total effects on city extent and population are ambiguous.

Gupta et al. [2021] look at how the housing market changed during the first year of Covid. To do this, they assemble data describing real estate transactions and their distance from center of the city. For rental and sales prices they rely on Zillow price indexes. These indexes are available at the zipcode-month level from the real estate website of the same name, and are intended to describe the price or rent for a "standard" house.

Gupta et al. [2021] use these data to estimate house price and rental price gradients for US cities, just as we already saw done for two Japanese and French cities in figure 1.1. The difference is that they use US data, and they estimate gradients for rental prices and for sale prices.

Letting V be sales prices, R be rental price, x be distance to the CBD, and letting i index zipcodes, Gupta et al. [2021] estimate

$$\ln V_i = A_1 + B_1 \ln x_i + \epsilon_i,$$

and

$$\ln R_i = A_2 + B_2 \ln x_i + \mu_i,$$

pooling data for the 30 largest MSAs in the US. This gives them a sort of average rent gradient for the largest cities in the US.

They conduct both regressions twice, once before the pandemic, and once after. Figure 2.4 reports their results. The left panel reports rent gradients with the logarithm of distance to the CBD on the x-axis and the logarithm of rent on the y-axis. The green line reports the pre-Covid rent gradient and the red line the post-Covid rent gradient.

As we expect, (1) the pre-Covid rent gradient is downward sloping, (2) the post-Covid, the rent gradient is flatter, and (3) post Covid, the intercept of the rent gradient is lower. Surprisingly, the rent gradient is actually slightly upward sloping post-Covid. This could reflect one of four things. First, that the Covid (dis)amenity so strongly favored suburban locations that it reversed slope of the rent gradient. Second, these rent gradients are actually averages over a sample of 30 MSAs, and that this averaging is partly responsible for the changes we see.[1] Third, these figures

[1] For example, if suburban rents increased more in larger cities, which are more heavily represented in the sample of "big x" rents.

Figure 2.4: Changes in real estate prices and rents from the Covid pandemic

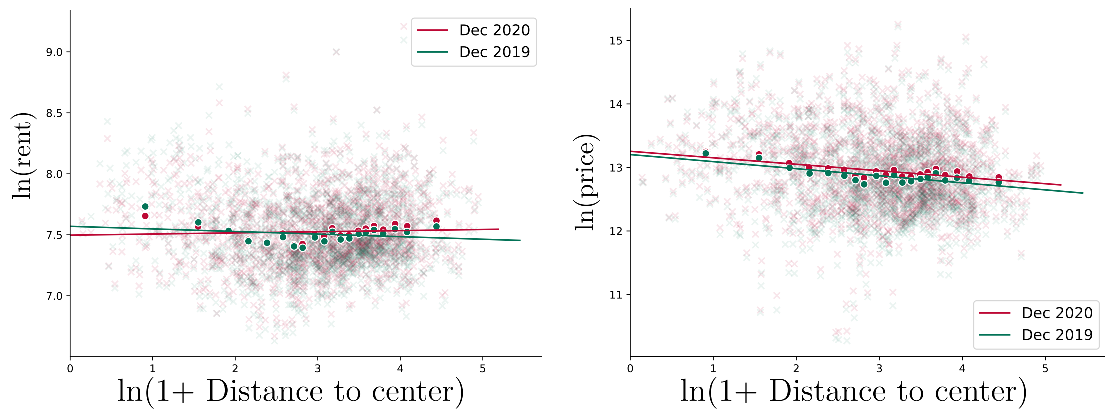

Note: *The left panel shows the relationship between log distance from the city center and log rent before (green) and after (red) the pandemic. The right panel is identical but reports sale price gradients. Lighter points indicate ZIP codes, while colored points indicate averages by 5% distance bins. Figure reproduced from Gupta et al. [2021].*

are based on a "Zillow rent index", which is supposed to be the rent for a standard house. This could be a problem if the "standard" rental unit in a suburban location changed from a small apartment before the pandemic to a large house during, but no corresponding change happened in central locations. Fourth, there is some serious problem with the monocentric city model, and this upward sloping rent gradient is alerting us to this problem.

The right panel of figure 2.4 is identical to the left, except that it reports gradients for sale prices, rather than rent. Here we see that both gradients are downward sloping and that the post-Covid gradient is flatter than the pre-Covid gradient. We do not see a decrease in the intercept of the post-Covid price gradient however. This probably reflects an increase in the overall demand for housing during the pandemic. This suggests an important omission from the monocentric city model, and one that we will address in Chapter 3. The monocentric city model does not allow people to choose their housing consumption. Everyone must choose housing, really land, $\bar{\ell}$.

Notice that average rents move more than average sales prices, and in particular, central rents decrease dramatically while central sale prices increase. Much was written during the pandemic about the possibility that Covid would lead to the end of cities as we know them. Some feared that as people retreated to remote work suburbs the central cities would remain only as blighted husks of their former selves.

Recall our discussion of how rental prices and asset prices are related in section 1.5.2. Asset prices are the discounted present value of rental prices. With this in mind,

Figure 2.5: Development patterns around Austin, Texas.

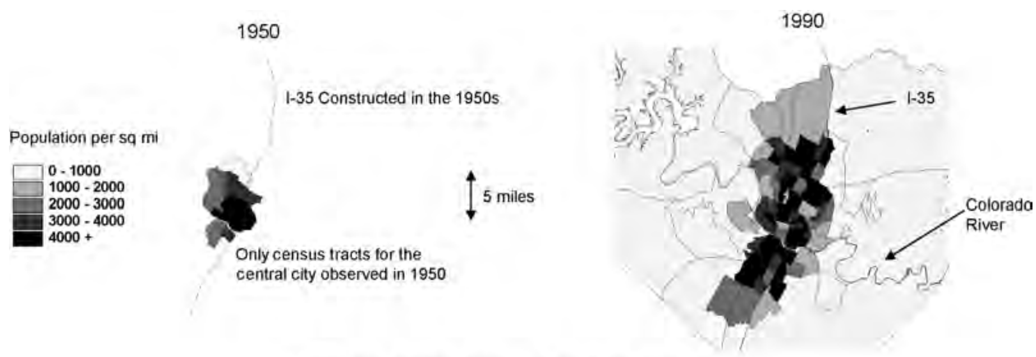

Note: *Population per square mile in census tracts near central Austin in 1950 and 1990. Between 1950 and 1990, Interstate 35 was constructed more-or-less north to south through central Austin. Over time, the city decentralized along the highway. Figure reproduced from Baum-Snow [2007].*

what does figure 2.4 suggest about the permanence of Covid related changes to the rent gradient? If people anticipated that the Covid related changes were permanent, then the rent and price gradients would change in exactly the same way. That the rent gradient declines in the center, even as asset prices increase slightly, must mean that the buyers of those central houses expect central cities to recover at least some of their pre-Covid attractiveness. Thus, pandemic fears about the death of cities were not widely shared by people who were actually deciding where to live.

Summing up, the monocentric city model seems to do pretty well at predicting the way that rent and price gradients respond to the pandemic decrease in transportation costs and change in central versus suburban amenities. Both rental and price gradients flattened out as we expect, and rental prices fell in the center. Two features of price and rent changes are noteworthy. First, we observe an upward sloping rent gradient post-Covid. This not strictly inconsistent with the model, if the changes in amenities are large enough it could reverse the slope of the rent gradient, but it is suspicious. Second, asset prices don't fall in the CBD but rental prices do. This probably the expectation that rents will rise in the future after the pandemic has passed.

2.3 Highways and decentralization

The US Interstate Highway System is a system of limited access highways built where no roads previously existed, or as upgrades of smaller highways. Construction of the Interstate began around 1955 with most of the network was complete by 1970. Most construction since 1970 has been of expansion lanes on existing routes. Not many other changes to US infrastructure more clearly reduced the cost of travel in general, and commuting in particular. We here consider how these new highways changed US cities, and whether these changes are consistent with what the monocentric city model predicts should happen when commuting costs fall.

In a classic paper, Baum-Snow [2007] investigates how US cities changed between 1950 and 1990 in response to the new highway network. He starts by constructing constant boundary MSAs for 139 MSAs. For each MSA he also constructs constant boundary central cities for these MSAs. These are the "old downtowns" as of 1950. He then asks what happened to population in these old downtowns during the period from 1950-1990 when the Interstate was constructed.

He finds that between 1950 and 1990, the total population of the 139 MSAs in his sample increased by 72%. Over the same period, the population of the old center cities *decreased* by 17%. That is, as the population of US MSAs nearly doubled, the population of their centers fell by almost one fifth. This sort of spreading out and growth is exactly what we see in the moncentric city model when commuting cost t falls.

Figure 2.5 illustrates the changes that took place in Austin, Texas around the time that Interstate 35 was constructed running North to South through the center of the city. Between 1950 and 1990, the city expanded linearly along the highway. Although Austin is not really a linear city, the changes illustrated in this figure look like the comparative static that we found in section 1.3.3, people spread out along the road as the cost of traveling on the road falls.

As a way to check whether this phenomena is general, Baum-Snow [2007] estimates a population density gradient, much like the land rent gradients we've already seen. To understand this regression, introduce the following notation. Let P_{ij} be population in census tract j of MSA i. Let $\text{dis}_{ij}^{\text{cbd}}$ be the distance from this tract to the CBD, and let $\text{dis}_{ij}^{\text{hwy}}$ be the distance to nearest interstate. Baum-Snow estimates,

$$\ln P_{ij} = \alpha_i + \beta \ln \text{dis}_{ij}^{\text{cbd}} + \gamma \ln \text{dis}_{ij}^{\text{hwy}} + \epsilon_{ij} \tag{2.1}$$

using data for 1950, and again for 1990.

This is a population density gradient, but elaborated to allow for the fact that transportation costs are not the same in every direction traveling outward from the CBD. Transportation costs ought to be lower along a highway than not.

Equation (2.1) is a population density gradient. As described in Chapter 1, the monocentric city model does not allow for variation in population density. Every household lives on a parcel of size $\bar{\ell}$. Thus, mechanically, equation (2.1) seems unlikely to help us evaluate the monocentric city model. This is a fair point. However, equation (2.1) will let us look at how people spread out from the center when transportation costs fall, a prediction of the model.

Baum-Snow finds that β increases from $-.132$ to $-.114$ between 1950 and 1990. Because the regression is conducted in logarithms, we can interpret the coefficients as elasticities. Thus a 1% increase in distance to the CBD leads to a 0.13% decrease in density in 1950, but only a 0.11% decrease in 1990. That is, the population density gradient got flatter. This is just what the monocentric city model predicts will happen when transportation costs fall. Baum-Snow also finds γ decreases from about -0.014 to -0.019, so population density falls faster as we travel away from a highway in 1990 than in 1950. This is not something the monocentric city model can address. The model assumes that the city is radially symmetric, and this regression is looking for changes that are not radially symmetric.

There is a problem with these estimates of equation (2.1), however. We cannot be sure whether changes in population density occur because people change their choice of residential location in response to the change in transportation costs, or whether the Federal Highway Administration cleverly anticipated where people were moving and built roads to meet these changes. In all likelihood, some of both is going on, and this means that interpreting equation (2.1) is difficult.

To resolve this problem, we would like to do something in the spirit of a clinical trial. That is, we imagine assigning each MSAs to a "treated" or "control" sample on the basis of a coin toss (heads and Brownsville is "Treated", tails and Spokane is "Control") and then assigning highways to the treated, but not the control sample in 1950. In 1990, we could learn the effects of highways on the treated cities by comparing them to the control cities. If we were able to do this, then we could be pretty sure that the treated cities were different from the control cities, on average, only because they got highways.

Simply describing this process, makes it clear that it's impractical. However, one of the important innovations in economics over the past 20-30 years has been the development of econometric techniques that allow us to simulate, sometimes more convincingly than others, exactly this sort of experiment. This is just what Baum-Snow does to resolve the problem of reverse causation, and much of his paper is about these econometric details. This allows Baum-Snow to estimate how much cities change in response to the construction of the interstate highway network.

Baum-Snow's econometric method requires that he restrict himself to a particular question. Define a "radial interstate ray" to be just what it sounds like, an interstate highway that travels from the constant boundary center city out of the city. In figure 2.5, I35 counts as two rays, one going north and one south. Let rays_i denote the count of radial interstate rays in city i, e.g., two for Austin, and let N_i^c be the center city population. The operator Δ indicates changes from 1950-90, and "controls" is a list of other variables whose purpose is to solve econometric problems beyond the scope of this book. Baum-Snow's main estimating equation is

$$\Delta \ln N_i^c = \delta_1 + \delta_2 \text{rays}_i + \text{controls}_i + \epsilon_i. \tag{2.2}$$

Because there were no interstate rays in 1950, the count of rays in 1990 is also the change in the number of rays, so this is really a regression of the change in the logarithm of central city population on the change in interstate rays.

To understand what this equation is doing, you need to recall some of the rules for manipulating logarithms,

$$\begin{aligned}\Delta \ln N_i^c &= \ln N_{1990i}^c - \ln N_{1950i}^c \\ &= \ln \frac{N_{1990i}^c}{N_{1950i}^c} \\ &= \ln(1 + r_i) \\ &\approx r_i\end{aligned}$$

So, as long as the rate of change is small enough that $\ln(1 + x) \approx x$ is a good approximation, the coefficient of rays in equation (2.2), δ_2, tells us the change in central city population caused by each ray as a share of the initial value.

Baum-Snow's big result is that δ_2 is about -0.11. This means that each radial interstate ray reduces central city population by about 11%. Because an average MSA received about 1.5 radial interstate rays, for an average MSA, the interstate caused about a 16% decline in the population of the constant boundary central city between 1950 and 1990. Recalling that the population of constant boundary central cities declined by about 17% during this time, this means that the interstate caused almost the entire decline in central city population between 1950 and 1990.

If we think that the main effect of highways on cities is to reduce transportation costs, then this looks pretty good for the monocentric city model. The monocentric city model predicts a decreasing share of population near the center as t falls, just what Baum-Snow finds.

2.4 Highways and growth

A second prediction of the (open) monocentric city model is that cities will grow when transportation cost falls. Duranton and Turner [2012] examine this hypothesis by looking at how MSA population (really employment) changes with lane kilometers of interstate highway between 1983 and 2003.

In their sample of 227 MSAs, average employment grew from about 250 thousand to about 410 thousand between 1983 and 2003, an increase of about 65%. During this same time, kilometers of interstate highway in an average MSA increased from

234 to 255, an increase of about 9%. This works out to annual growth rates of about 2.8% for population, and 0.5% for highway kilometers.[2]

We hope that the Federal Highway Administration builds highways in cities where people want to move, so we should be concerned that correlations between changes in MSA highways and changes in employment could reflect reverse causation. Duranton and Turner rely on the same basic econometric technique as Baum-Snow [2007] to try to overcome this problem, and like Baum-Snow [2007], much of the Duranton and Turner paper is devoted to describing this technique. As I did in the discussion of Baum-Snow, I'm going to skip over these details.

Duranton and Turner want to check whether cities with more highways grow faster. For this purpose, let n_{it} be employment in MSA i at year t, let r_{it} be lane kilometers of interstate in MSA i in year t, and let x_{it} be control variables that address econometric problems beyond the scope of this book. The main estimating equation from Duranton and Turner [2012] is,

$$\Delta \ln n_{it+1} = A_0 + A_1 \ln r_{it} + A_2 \ln n_{it} + A_3 x_{it} + \varepsilon_{it} \tag{2.3}$$

This looks a lot like the Baum-Snow regression, equation (2.2), but there is an important difference. Because Baum-Snow started his study when there were zero interstates, his control for highway rays was really "change in rays". That's not what's happening here. Here employment growth is a function of the initial level of highway lane kilometers, so it's actually not easy to compare the two regressions, even though they look a lot alike.

As before,

$$\begin{aligned}\Delta \ln n_{it+1} &= \ln n_{it+1} - \ln n_{it} \\ &= \ln(n_{it+1}/n_{it}) \\ &= \ln(1 + \rho_n) \\ &\approx \rho_n\end{aligned}$$

So that A_1 tells us the effect on the growth rate of the MSA employment, ρ_n from a change in initial lane kilometers of interstate.

The main empirical result in Duranton and Turner is that A_1 is about 0.15. This means that a 1% increase in lane kilometers increases the annual employment growth rate for an average MSA by about $0.15 \times 1\% = 0.15\%$. Recalling that the stock of interstate kilometers in an average MSA grows by about 0.5% per year during their

[2]If you are paying close attention, you will notice that the number of MSAs varies across the different studies I've described. This reflects three things: (1) sometimes, not all data is available for all MSAs, (2) the number of MSAs increases over time as more metropolitan areas cross the 50,000 population threshold, and (3) the studies may rely on slightly different definitions of MSA.

Figure 2.6: Average city population growth rate as time to subway system opening varies

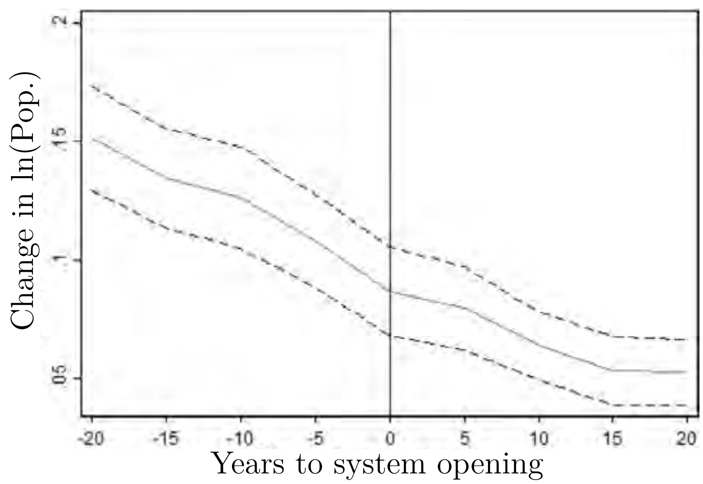

Note: *Subway system opening and population growth (constant sample of 61 cities). The graph depicts mean change in city log population according to time to system opening.* $t = 0$ *indicates the year in which a city's subway system was inaugurated. We impose a constant sample of cities on either side of* $t = 0$*. Graph based on constant sample of 61 cities. It does not look like subways cause population growth. Figure reproduced from Gonzalez-Navarro and Turner [2018].*

sample, highway construction contributes $0.15 \times 0.5\% = 0.075\%$ per year to the baseline 2.8% annual growth rate of MSA employment. The monocentric city model predicts that a city will get bigger when transportation costs fall. This is just what Duranton and Turner find. As we build more roads in a city, more people come, but the magnitude is tiny.

2.5 Subways, decentralization and growth

If changes to transportation costs affect the way cities are organized, it shouldn't matter if changes result from better roads, more telecommuting, or better subways. We've just checked telecommuting and roads.

Gonzalez-Navarro and Turner [2018] check subways. Their paper is based on three main sources of data. The first is a census of all subway systems in the world. The second is a panel of data describing the population of all 632 cities in the world that had a population above 750,000 sometime between 1950 and 2010. The third is "lights at night" data for these cities. This is the same data we saw in figures 2.1 and 2.2.

Figure 2.7: Lights at night and subways in Mexico City and Beijing in 2010

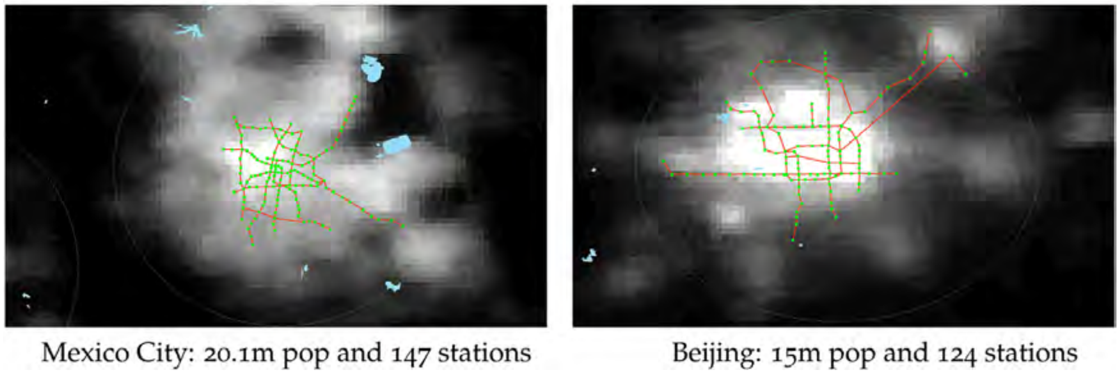

Note: *2010 lights at night, subway route maps, and all subway stations constructed prior to 2010 in Mexico City (left) and Beijing (right). The gray/green ellipses in each figure are projected 5 km and 25 km radius circles to show scale and light blue is water. Figure reproduced from Gonzalez-Navarro and Turner [2018].*

As of 2010, among these 632 large cities there are 138 with subway systems. Four subways were in operation around 1860, Liverpool, Boston, London and New York. Cities in Asia began constructing subways in the 1970 and account for most of the new systems since that time. Overall, cities in Europe are much better provided with subways than anywhere else. As of 2010, an average system consists of 77 route kilometers and 57 stations. On average, the 138 subway cities had a population of about 4.7m people in 2010 and their populations grew at an average rate of just above 2% per year between 1950 and 2010.

Among the 138 subway cities that Gonzalez-Navarro and Turner [2018] study, only 61 have data that is complete enough to allow the calculation of population growth rates 20 years before and 20 years after the system opening. The solid line in figure 2.6 reports the average population growth rate in these cities as a function of the time to the subway system opening. The dashed lines report confidence intervals around the mean. That this line slopes downward tells that on average the growth rate is falling in these cities. This is a common empirical finding. Researchers often find that the growth rate of cities falls as they get larger, and this sample consists of large cities, getting larger. More interestingly, we see no change in the trend around the time when each city's subway system opened (zero on the x-axis). This is one of the main findings from Gonzalez-Navarro and Turner [2018]; the opening of a subway system does not seem to affect the level or growth of population in the cities where they open.

Gonzalez-Navarro and Turner also investigate whether subways decentralize cities.

To do this, they use lights at night data to estimate a light gradient and then ask whether this gradient flattens after the subway system opens.

This means estimating more gradients, this time to see how the intensity of lights at night varies with distance to the center. For each of 138 subway cities, for each year when they observe night lights (1995, 2000, 2005, 2010), they calculate mean light intensity in a series of donuts, 0-1.5km, 1.5-5km, 5-10km, 10-25km and 25-50km, centered on the CBD. The 5km and 25km donuts are faintly visible in figure 2.7. Next, let y_{itd} be the mean light intensity in donut d for city i in year t, and let x_{itd} be distance of the midpoint of the donut from center, e.g., 7.5 km for 5-10km donut. They can now estimate city-year specific light density gradients,

$$\ln y_{itd} = A_{it} + B_{it} \ln x_{itd} + \epsilon_{itd}.$$

That is, Gonzalez-Navarro and Turner estimate the slope of the light gradient 138 times in each of the four years they observe lights at night. This gives a separate B_{it} for each city-year, each describing the change in mean light intensity as we move from the smallest to the largest of the five donuts. These are exactly the same log-linear gradients we've already seen, but explaining the intensity of lights at night rather than property prices or population density.

With the slope B_{it} for city-years it in hand, Gonzalez-Navarro and Turner can check if subways cause cities to spread out with the following regression,

$$\Delta B_{it} = A_0 + A_1 \Delta \ln(\text{Subway Stations}_{it}) + A_3 \text{Controls}_{it} + \epsilon_{it}$$

because light gradients are downward sloping, if subways cause cities to decentralize, they will increase the B_{it}, and this is just what they find. That is, as the number of subway stations increases, B increases and the light gradient gets flatter. This is just what the monocentric city model predicts will happen when transportation costs fall.

2.6 Amenities and city size

Another prediction of the monocentric city model is that cities will be bigger as their amenities are better. To check this, Glaeser and Gottlieb [2009] collect data on 316 US MSAs in 1990 and 2000, and ask how population growth is related to weather.

They find a strong relationship between good weather and growth in population. Figure 2.8 illustrates this result. On average, a one degree Fahrenheit increase in mean January temperature is associated with a 2% increase in population over the course of a decade. This is also consistent with the predictions of the monocentric city model.

Figure 2.8: Population growth and mean January temperature for US MSAs between 1990 and 2000

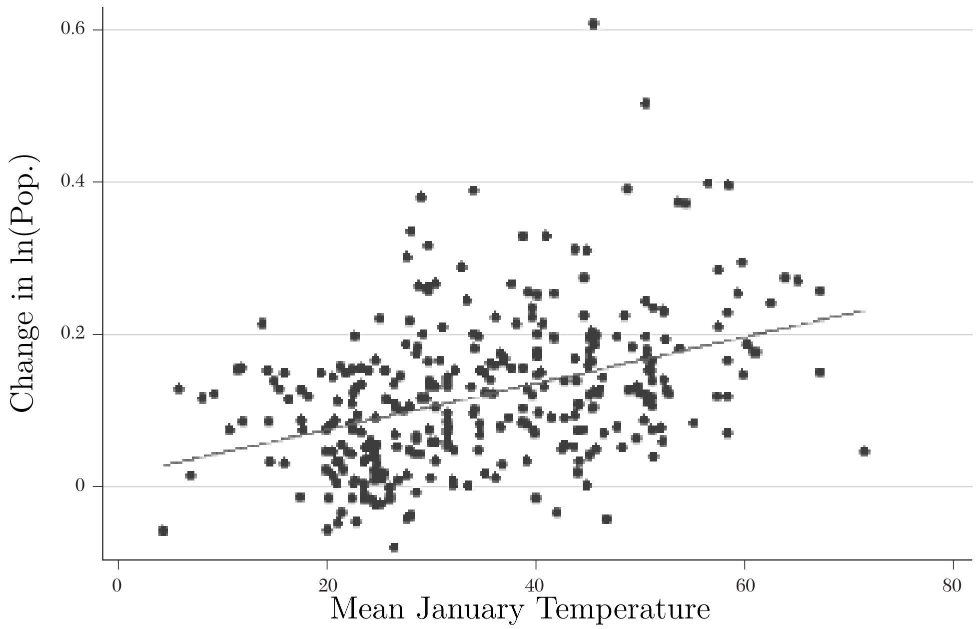

Note: *y-axis is the change in the logarithm of MSA population between 1990 and 2000. x-axis is mean January temperature. US MSAs with milder winters grow faster. Figure reproduced from Glaeser and Gottlieb [2009].*

2.7 Property taxes and land prices

One of the more interesting consequences of spatial equilibrium is that property taxes are capitalized into land prices in a really mechanical way. One dollar of property taxes equals one dollar of rent, and one dollar of property taxes per year equals the discounted present value of one dollar per year in asset prices.

In reality, things ought to be more complicated for two reasons. First, property taxes are assessed on the value of land *and house*. If you put an addition on your house, you need to pay property taxes on the value of the addition forever. Similarly for a new paint job, etc. Thus, if we allow a little more realistic description of the world, we might expect that 1$ of property taxes will decrease the value of house and land by more than 1$ because it will lead to sub-optimal maintenance.

Second, up to now, we have implicitly assumed that property taxes leave the model. They go to the city government and are entirely wasted. In fact, property taxes are used, at least in part, to provide important public services like trash collection, fire and police protection, parks and roads. These things will operate like amenities, and hopefully, have value of at least 1$ per dollar of taxes collected. Strictly, the monocentric city model predicts that property taxes decrease the value of a property *all else equal*, but it is not easy to find real world examples where taxes changes and services do not.

One case that seems pretty close occurred in 2008 in Toronto. In 2008, Toronto imposed a "land transfer tax". This is a property tax that you pay when you sell your property, rather than every year, as property taxes are usually collected. This tax was imposed in Toronto, but not in neighboring municipalities. Because the land transfer tax is attached to a municipality, we should expect their effect on real estate markets to vary discretely at municipal borders. If taxes are capitalized into real estate prices as the monocentric city model suggest, then we should see prices fall in Toronto by about the magnitude of the tax, net of whatever value of public services the tax will purchase. Dachis et al. [2012] argue that the circumstances surrounding the imposition of the 2008 Toronto land transfer tax make it unlikely that the new tax affects the provision of municipal services. They then do exactly the experiment described above. They find that real estate prices fall by about one dollar for every dollar of tax assessed (although their estimates are not very precise).

Palmon and Smith [1998] finds another way to look for changes in property tax rates, all else equal. They study data describing house prices in 50 subdivisions in the Houston suburbs. All 50 subdivisions are similar. They are served by three school districts of similar quality, and were constructed within a few years of each other.

Water and sewer service is the same for every subdivision and was provided by the private developers who built out each subdivision. Developers financed the construction of water and sewer infrastructure by issuing bonds, with the payment on these bonds financed by property taxes collected from homeowners in the subdivision. The interest rate on the bonds, and hence the subdivision property tax rate, varies with the interest rate that prevailed when construction occurred. This means that the different subdivisions are paying different prices for the same water and sewer service. They find that about 65 cents of every dollar of property tax is capitalized into property prices. This suggests that the basic logic of capitalization that emerges from the monocentric city model is economically important. But it is also a bit less than the predicted 100% capitalization, so the model, somehow, not exactly right. It's worth noting that this analysis is based on only about 500 real estate transactions, so we should worry about how precise their estimates are and whether they apply outside of the Houston suburbs.

Figure 2.9: Median expenditure on rent divided by median wage, in a sample of MSAs

Table 1
Median ratio of rental expenditures to wage and salary income, median income, and growth in real rental prices.

MSA	Median ratio			Median HH income (2000) renters only	Real rent growth, 1980–2000
	1980	1990	2000		
Albany–Schenectady–Troy	0.21	0.23	0.23	$32,300	16.2%
Atlanta–Sandy Springs–Marietta	0.24	0.25	0.25	$36,300	25.1%
Austin–Round Rock	0.27	0.25	0.25	$36,400	42.0%
Bakersfield	0.28	0.25	0.25	$26,800	0.7%
Baltimore–Towson	0.23	0.23	0.23	$34,000	35.1%
Boston–Cambridge–Quincy	0.24	0.26	0.24	$43,000	52.1%
Buffalo–Niagara Falls	0.20	0.22	0.23	$28,800	21.1%
Charlotte–Gastonia–Concord	0.23	0.24	0.24	$37,000	27.3%
Chicago–Naperville–Joliet	0.21	0.23	0.23	$36,000	33.5%
Cincinnati–Middletown	0.21	0.22	0.20	$30,400	5.5%
Cleveland–Elyria–Mentor	0.21	0.22	0.23	$30,000	5.1%
Columbus	0.22	0.23	0.23	$33,100	38.6%
Dallas–Fort Worth–Arlington	0.24	0.24	0.24	$34,600	32.5%
Denver–Aurora	0.25	0.24	0.26	$35,000	19.5%
Detroit–Warren–Livonia	0.21	0.22	0.22	$35,000	6.8%
Fresno	0.25	0.27	0.26	$25,900	14.0%
Grand Rapids–Wyoming	0.19	0.24	0.21	$31,000	16.9%
Greensboro–High Point	0.24	0.23	0.22	$33,000	23.7%
Houston–Sugar Land–Baytown	0.23	0.22	0.23	$32,000	7.2%
Indianapolis–Carmel	0.21	0.23	0.23	$34,000	8.6%
Jacksonville	0.27	0.24	0.25	$31,000	3.5%
Kansas City	0.21	0.22	0.22	$35,700	21.4%
Las Vegas–Paradise	0.29	0.27	0.27	$35,000	20.1%
Los Angeles–Long Beach–Santa Ana	0.25	0.29	0.27	$33,000	37.2%
Louisville–Jefferson County	0.22	0.23	0.21	$32,000	4.2%
Miami–Fort Lauderdale–Pompano Beach	0.27	0.29	0.29	$28,000	24.3%
Milwaukee–Waukesha–West Allis	0.20	0.23	0.22	$32,000	12.2%
Minneapolis–St. Paul–Bloomington	0.24	0.25	0.23	$35,500	19.3%
Nashville–Davidson–Murfreesboro–Franklin	0.23	0.24	0.24	$33,000	22.9%
New Orleans–Metairie–Kenner	0.24	0.25	0.24	$25,000	24.6%
New York–Northern New Jersey–Long Island	0.22	0.24	0.24	$39,600	38.2%
Orlando–Kissimmee	0.26	0.27	0.27	$32,950	41.1%
Philadelphia–Camden–Wilmington	0.22	0.24	0.23	$37,000	33.2%
Phoenix–Mesa–Scottsdale	0.28	0.26	0.26	$32,000	9.5%
Pittsburgh	0.21	0.21	0.22	$30,000	10.1%
Portland–Vancouver–Beaverton	0.27	0.24	0.25	$36,000	19.1%
Riverside–San Bernardino–Ontario	0.26	0.28	0.27	$32,000	17.9%
Sacramento–Arden-Arcade–Roseville	0.25	0.28	0.26	$33,000	38.9%
St. Louis	0.22	0.23	0.22	$30,000	4.4%
Salt Lake City	0.24	0.23	0.27	$30,900	22.5%
San Antonio	0.22	0.24	0.24	$30,000	13.5%
San Diego–Carlsbad–San Marcos	0.29	0.30	0.28	$34,000	38.4%
San Francisco–Oakland–Fremont	0.26	0.28	0.25	$46,900	70.7%
San Jose–Sunnyvale–Santa Clara	0.24	0.26	0.25	$58,500	110.0%
Seattle–Tacoma–Bellevue	0.25	0.25	0.26	$38,200	33.7%
Syracuse	0.24	0.24	0.24	$27,000	16.7%
Tampa–St. Petersburg–Clearwater	0.26	0.25	0.25	$31,400	23.0%
Tucson	0.26	0.29	0.26	$24,600	−2.7%
Tulsa	0.23	0.22	0.23	$31,000	1.8%
Washington–Arlington–Alexandria	0.23	0.26	0.24	$44,600	46.7%
Average	0.24	0.25	0.24	$33,689	24.2%
Standard deviation	0.02	0.02	0.02	$5710	19.4%

Note: *Table reproduced from Davis and Ortalo-Magné [2011].*

2.8 Wages and rents

Another prediction of the monocentric city model, and the last one we'll check, is that a 1$ increase in wages leads to a 1$ increase in land rent.

We can find some evidence about this in Davis and Ortalo-Magné [2011]. This paper looks at the relationship between income and expenditure on housing using a large census data set. Table 2.9 presents their findings. For each MSA and census year, they calculate the ratio of rent to income for each household. The table presents the value of this ratio for the median household in each MSA. The last row of the table reports the average of these values over all MSAs. That is, the average median rent to income ratio. The median share of income devoted to rent ranges between about 0.21 and 0.29, with most MSAs even closer to 0.25. That is, people spend about 25% of their income on housing, no matter where they live (in the US) or how much money they make.

Glaeser and Gottlieb [2009] also make this point. In particular, for MSAs in the US in 2006, they regress the logarithm of median MSA income on the logarithm of median home value. That is,

$$\log(\text{MSA median income})_i = A_0 + A_1 \log(\text{MSA median home value})_i + \epsilon_i.$$

They find that $A_1 = 0.34$. This means that a 1% increase in MSA mean income is associated with a 0.34% increase in MSA mean home value. This is a little larger then the 0.25 that Davis and Ortalo-Magné [2011] finds, but the data are a little different too.

Clearly wages and housing expenditure move together as the monocentric city model requires. But, equally clearly, the effect is much less than the one-for-one relationship that the model predicts.

2.9 Conclusion, the state of the evidence

The first prediction of the monocentric city model is that the rent gradient, $R(x)$, should decrease with distance to the center. This is broadly consistent with observation. We see it for cities in France and Japan in figure 1.1, for housing price gradients in the US both before and post-Covid, and for apartment rental prices before Covid in figure 2.4. The post-Covid rental price gradient for apartments increases slightly as we move away from the center. This contradicts the prediction of the simplest version of the monocentric city model, though if we allow residential locations to contribute directly to utility, i.e., "amenities", then this slight upward slope could reflect capitalization of Covid risk in a way that is consistent with the model.

The next prediction of the monocentric city model is that rent gradients should flatten as commuting costs fall. This is exactly what we saw during Covid, commuting costs fell as people began to work remotely. In figure 2.4 we see that both sale and rental price gradients flattened in response, although part of this response was surely due to the greater risk of Covid in denser more central locations.

The monocentric city model predicts that cities will spread out and grow in size as transportation costs fall. Between 1950 and 1990, the population of constant boundary center cities in the US fell by 17%, even as the population of MSAs that contained them increased by 72%. The entire decrease in central city population can be accounted for by the construction of Interstate highway rays through these cities. If we think that the main effect of highways was to reduce transportation costs, then this is broadly consistent with the prediction of the monocentric city model. Gonzalez-Navarro and Turner [2018] finds something similar for subways. The gradient of light intensity flattens out in cities after they build subway systems.

We need a caveat here. Baum-Snow [2007] finds that the population of the constant boundary central city falls with the construction of radial interstate highways. This means that population density in the center must fall. As we formulated the monocentric city model in Chapter 1, each household is constrained to consume a constant amount of land $\bar{\ell}$, and so population density cannot change. This is an obvious shortcoming of the model, and we address it in Chapter 3.

The monocentric city model predicts that the population of a city will grow with a reduction in commute costs. Here, the evidence is for a small effect, at most. Gonzalez-Navarro and Turner [2018] can find no effect on city populations from the construction of subways, and Duranton and Turner [2012] find a small effect of interstate highways on the growth of employment. There is nothing in the monocentric city model that requires that reductions in commute costs have a large effect on population, but it does require it to be positive. The available evidence is not conclusive on this point. The effects of transportation infrastructure on city population are small or zero, and it is not clear which.

If wages, w, increase by one dollar, the monocentric city model predicts that expenditure on housing will also increase by one dollar for almost all households. In fact, in the US expenditure on housing increases by about one dollar for each three or four dollars of additional income. This is the right sign, but the magnitude is too small. However, this comparison is not quite fair. The data describes expenditure on "housing" that consists of both house and land, while in the model "housing" is just land. Just as we need a richer model of housing in order to think about population density, we also need a richer model of housing if we are going to predict the share of wages that are capitalized into housing prices.

Property taxes are also capitalized into asset price much in the way the model suggests. Because the model requires that tax revenue leave the model, when in reality at least some of it will provide valuable services, the model is addressing an unusual special case. In the rare real world analogs of this special case, a good guess would be that 60-100% of a household's property tax bill is capitalized into the price of housing. That is, if the property tax on a house goes up by one dollar per year forever, then we expect the sale price of the house to fall by 60-100% of the discounted present value of this stream of payments.

If we generalize the basic model to allow for locations to contribute directly to utility through amenities, then the model predicts that cities will be larger as amenities improve. We have evidence in support of this. US cities with milder winters grow faster than those with harsher winters.

We can now score the contest between the monocentric city model and the data. Almost all of the predictions the model makes about the slope of the rent gradient and about the way that a city responds to changes in commuting costs, wages, and amenities are qualitatively correct. Some of the details and magnitudes are wrong, however. Rent goes up only 25 to 35 cents for every one dollar increase in wages, not one dollar as the model predicts. Central city population density falls as commuting costs fall but the model does not allow population density to change. The model predicts that land rent fall about linearly with distance to the center, while the data shows a much faster decrease.

Finally, the model predicts that cities grow as commute costs fall, and in spite of the fact that the data shows that changes in commute costs have large effects on the way cities are organized, we see at most a tiny effect of changes in commute costs on city size. Given how successful the monocentric model is otherwise, this finding is puzzling.

Problems

This problem will examine the change in the rent and purchase price gradients from Gupta et al. (2021).

1. Before the pandemic, the rental price gradient was described by:

$$\ln R_0(x) = 7.6 - 0.04 \ln(x + 1)$$

   where x is distance from the city center. This is shown in panel A of Figure 3 from Gupta et al. (2021). During the pandemic, the rental gradient changed to:

$$\ln R_1(x) = 7.5 - 0.004 \ln(x + 1)$$

   What are the monthly rental prices at $x = 0$, before and during the pandemic? What is the percent change in rent at $x = 0$?

2. As shown in Panel B, the asset price gradient before the pandemic was described by:

$$\ln P_0(x) = 13.2 - 0.127 \ln(x + 1)$$

   During the pandemic, this gradient changed to:

$$\ln P_1(x) = 13.15 - 0.115 \ln(x + 1)$$

What are the asset prices at $x = 0$, before and during the pandemic? What is the percent change in asset price at $x = 0$?

3. Suppose that the pandemic-related changes in rental prices are permanent. Use 1 to find the implied asset price of rental properties at $x = 0$ before and after the pandemic, using interest rate $\rho = 0.03$. What is the percent change in these implied asset prices?

4. Compare this implied change in asset prices, which assumed that the change in rental prices due to Covid would be permanent, to the actual change in asset prices from 2. Which is larger? What does this suggest about how long people expect the pandemic to last?

5. Throughout the pandemic, people have speculated that Covid would be "the death of cities". What does your work above suggest about this sort of speculation?