

A Field Guide to Urban Economics

Matthew Turner

V1 – Compiled April 4, 2025

Copyright Matthew Turner, 2024. All rights reserved. No part of this document may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the author

Contents

1	The Monocentric City Model	1
1.1	The Realtor’s mantra and spatial equilibrium	1
1.2	Land rent gradients in real life	4
1.3	The Monocentric City Model	7
1.3.1	The monocentric city model in two pictures	11
1.3.2	Three extensions	13
1.3.3	Comparative statics	17
1.3.4	Land rent and welfare	22
1.4	Application #1: Learning the value of school quality from real estate prices	24
1.5	Application #2: Detecting racist property tax assessments	26
1.5.1	Property taxes and rental prices	28
1.5.2	Land rent and capitalization	31
1.5.3	Fair assessment of property taxes	31
1.6	Conclusion	33
2	The Monocentric City Model vs. Data	37
2.1	Cities in real life	37
2.2	Rent gradients and Covid	40
2.3	Highways and decentralization	44
2.4	Highways and growth	48
2.5	Subways, decentralization and growth	49
2.6	Amenities and city size	52
2.7	Property taxes and land prices	53
2.8	Wages and rents	54
2.9	Conclusion, the state of the evidence	56
3	The Monocentric City Model with Housing	59
3.1	Population density gradients in real life	59
3.2	Households and housing	61

3.3	The construction sector	69
3.4	Summary	75
4	Urbanization in the Developed (mostly US) and Developing World	79
4.1	Urbanization in the US	80
4.2	Urban productivity premium	82
4.3	Excess urban mortality	85
4.4	The Monocentric City Model and urbanization in the developed world	88
4.5	Urbanization in the developing world	92
4.6	The Monocentric City Model and Rural Amenities	105
4.7	Conclusion	112
5	White Flight, Gentrification, and Bid-Rent	115
5.1	Transportation costs and decentralization	117
5.2	The Great Migration and white flight	119
5.2.1	The American Ghetto, 1890-1990	120
5.2.2	Causes of ghettos	124
5.3	Transportation costs, White Flight, and decentralization	130
5.4	Sorting and bid-rent	130
5.5	Conclusion	140
6	Quantitative Spatial Models and How Railroads Reorganized London	143
6.1	Distance in discrete space	145
6.2	The discrete choice problem	148
6.3	Extreme value distributions	153
6.4	A discrete linear city with heterogeneous households	156
6.5	Welfare	159
6.6	A discrete city with heterogeneous households and iceberg commute costs	162
6.7	The big prize: choosing discrete workplace <i>and</i> residence	163
6.8	Railroads, subways, and the economic geography of London, 1866-1921	167
6.9	Conclusion	174
7	The Roback Model and the value of amenities	178
7.1	What is the value of amenities?	178
7.2	Amenities and productivity in the moncentric city model.	179
7.3	Roback Model	181
7.4	Roback Theorem	188
7.5	Comments	192

7.6	Application #1: Valuing climate	192
7.7	Application #2: Public finance in spatial equilibrium	195
7.8	Conclusion	197
8	Agglomeration Economies, or why are there Cities Anyway?	200
8.1	Why are there cities?	200
8.2	Returns to scale and a planner's problem	201
8.3	Returns to scale and equilibrium	203
8.4	How to measure agglomerations?	207
8.5	How to measure agglomeration economies	212
8.6	Measuring agglomeration economies, empirical results	213
8.6.1	Spatial scale	215
8.7	Mechanisms	217
8.8	(Some of) what we know about mechanisms	218
8.8.1	Input sharing	218
8.8.2	Sorting and learning	219
8.8.3	Natural advantage and path dependence	219
8.8.4	Knowledge spillovers	222
8.9	Conclusion	223
9	Systems of Cities	227
9.1	Stylized facts about cities	227
9.2	Zipf's law	230
9.3	Systems of cities	235
9.4	Conclusion	241
10	Sorting, a Simple Hedonic Model, and Local Public Finance	243
10.1	Sorting vs causation?	243
10.2	Sprawl and obesity	244
10.3	Neighborhood and diet	245
10.4	Test scores and house prices	247
10.5	Neighborhood and human capital	248
10.6	Spatial equilibrium and Tiebout sorting	252
10.7	The Hedonic Model	256
10.8	Empirical evidence and conclusion	258

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 a bit like a punchline, but it is still important. First, it is an important observation about how the world works. Second, it highlights what is special about urban economics. Unlike most of the rest of economics, you can't study how cities are organized without thinking about locations.

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.

¹For 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 the 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 the house B to equal to 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 houses 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 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. In math jargon, 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 or, in urban economics jargon, 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, Hiratsuka and Yokohama, in 1991. 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 really 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 very different looking pairs of graphs are actually plotting the same information, but represented differently.

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 leaves variables in levels.

Eye-balling the figure for Paris, we see that the intercept is about 10 and the slope about $-\frac{1}{4}$. Writing this out, we have

$$\ln R = 10 - \frac{1}{4} \ln x. \quad (1.5)$$

We can rearrange this to get

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

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,

$$\ln n = A + B \ln r. \quad (1.7)$$

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

$$\ln n^0 = A + B \ln r^0 \quad (1.8)$$

and

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

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 jargon, B is the elasticity of n with respect to r .

1.3 The Monocentric City Model

We have 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, net of changes in rent, 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, $\bar{\ell}$, at the location x that they choose. Households use their wage to pay the land rent for their parcel, $R(x)\bar{\ell}$, 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

²This 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 \bar{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,

$$\begin{aligned} v(x) &= \max_x u(c) \\ \text{s.t. } w &= c + R(x)\bar{\ell} + 2t|x|. \end{aligned} \tag{1.10}$$

This should look familiar if you’ve studied microeconomics. It 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 ($u' > 0$) 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 \bar{u} . That is,

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

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 \bar{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 adjust so that all 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,

$$c^* = u^{-1}(\bar{u}). \quad (1.12)$$

For example, if $u(c) = c^{1/2}$ and $\bar{u} = 2$, 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,

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

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 \bar{R} . Let \bar{x} denote the distance of these boundary points from $x = 0$. At this location, we must have

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

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

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

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

We see here how the open city assumption works. A household must obtain a utility level of at least \bar{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 \bar{R} , so the price of the best unoccupied parcel is $\bar{R}\bar{\ell}$. 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 $-\bar{x}$ to \bar{x} and each household consumes an exogenously fixed amount of land, the population of the city is,

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

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

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

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

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

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

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. Given the assumptions that households optimize and that no one wants to move, 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, $u(c) = \ln(c)$, $\bar{R} = 0$, $\bar{u} = 0$ and $\bar{\ell} = 1$. Then, the household's problem is to choose location and consumption to solve,

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

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 $\ln(c) = \bar{u} = 0$ so that $c^* = 1$ everywhere.

Using $c^* = 1$ in the budget constraint, we have

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

which means that,

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

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

indicate units of consumption. Because w , c^* , \bar{R} and $\bar{\ell}$ 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, $\bar{R}\bar{\ell}$.


Looking again at the budget constraint, we have

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

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 $\bar{\ell}$ and to the cost of commuting.

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

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 red lines show how commute costs increase with distance from the CBD and the blue 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.*

exactly offsetting amount. This gives us the the blue 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

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.*

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 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, \bar{u} , we know the number of people who live in the city.

Given population size N , because everyone consumes $\bar{\ell}$ of land, the length of the

city must be $N\bar{\ell}$, and so the edges of the city must be at $|\bar{x}| = N\bar{\ell}/2$. 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,

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

Rearranging, we have $c^* = (w - \bar{R}\bar{\ell}) - 2t\bar{x} = w - (\bar{R} + tN)\bar{\ell}$. This means that rent at the CBD is $\bar{R} + tN$ per unit land, and so $c^* = w - (\bar{R} + tN)\bar{\ell}$ for the person at the center. With spatial equilibrium, everyone gets the same consumption. Thus, $\bar{u} = u(c^*) = u(w - (\bar{R} + tN)\bar{\ell})$.

With an open city equilibrium, the reservation utility is exogenous, and the population of the city adjusts until the marginal person pays just the agricultural rent. With a closed city equilibrium, population is fixed and utility adjusts so the person at zero pays just enough that the person at the edge doesn't want to outbid him for the central spot. Otherwise, everything else is the same.

Circular city

Suppose we relax the (silly) assumption that the city is on a line, and think about a symmetric circular city, still on a featureless plane, keeping everything else the same. This doesn't change the household's problem at all. We still have

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

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

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

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

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

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

Reorganizing, we have

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

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 $\pi\bar{x}^2$, so population is

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

This is the big 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.

Really, the circular city isn't much different from the linear city, except that you have to keep track of more math. So we'll stick with the linear city whenever we can.

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

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

Reorganizing, we have,


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

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 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.

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.

Figure 1.5: The monocentric city with amenities



Note: The blue line illustrates the land rent gradient for the baseline case when there are no amenities, really when $A = 1$. The green 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.

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.

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 how the monocentric city changes as we change wages or commuting costs, we've already looked at what happens when amenities change.

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 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

Box 1.3.2: Partial differentiation

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

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

This is the “instantaneous slope” of f at x .

Partial differentiation is the generalization of this idea to surfaces. Consider a function $F : R^2 \rightarrow R$, 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

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


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.

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 aside if you need help with this). To proceed, take the partial derivative of the land rent

Figure 1.6: Monocentric city comparative statics as commuting costs change



Note: The blue line describes an equilibrium land rent gradient in an open city when commuting costs are high, and the green 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.

gradient with respect to the commute cost,

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

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

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

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 $N = 2\bar{x}/\bar{\ell}$, we can use the chain rule to get

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

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: Monocentric city comparative statics as the wage changes



Note: The blue line shows an equilibrium land rent gradient for a low wage, and the green 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 is if 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.


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 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 unchanged. Balancing household budgets can only happen if most of 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

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



Note: The green line shows the equilibrium land rent gradient in a closed city with low commute costs, and the blue 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 away 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 $x = 0$.

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. As the commuting cost increases in a closed city, 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. Land rent at the edge of the city also stays constant because the rent required to bid land away from farmers also doesn't change. But in

equilibrium, the sum of land rent and commuting must be the same at all occupied locations. By assumption, the slope of the commute cost gradient is increasing. The slope of rent gradient has to increase to match. This means that the $x = 0$ intercept increases. It follows that land rent increases everywhere, except at \bar{x} where it is constant. Because income is fixed, consumption and utility have to fall, and total land rent increases. With a closed city, the costs of the increased commuting costs are divided between residents and landlords. This is quite different from the open city where the utility of households is fixed, and any change ends up affecting land rent without affecting household utility.


It is interesting to compare incidence of commuting changes in open and closed city. In an open city, the supply of people is perfectly elastic. All changes fall on landlords, good or bad. If you've studied microeconomics, recall the theory of tax incidence for ordinary goods. For ordinary goods, if demand is perfectly elastic, then the whole incidence of the tax falls on the seller. The intuition for this is the same as what we are seeing for the incidence of commuting cost changes in an open city. In an open city, the supply of people is perfectly elastic, and so changes in the value of living in the city fall on the landlord. 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 for policy evaluation of understanding how responsive is migration to local economic conditions.

1.3.4 Land rent and welfare

In an open city equilibrium, each household gets $u(c^*) = \bar{u}$, and they get this payoff no matter how much rent they pay. In this sense, land rent is a measure of the benefit to a household from living in the city. They can get payoff \bar{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 for this 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. If you've studied economics you will recognize that 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.

Figure 1.9: Aggregate land rent in the monocentric city



Note: The green line describes a land rent gradient for an open city. \bar{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, $y < \bar{x}$, then aggregate rent is less than for a city with edges at \bar{x} . In particular, the rent described by the two pink triangles is lost. If the planner chooses an extent greater than \bar{x} , then the planner must subsidize the marginal households to allow them to bid land away from farmers and still purchase c^* .

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, in equilibrium, the monocentric city is what happens when everyone pursues their own narrow self interest. What would happen if a rent maximizing developer organized the city, subject to free mobility for the households? Would the resulting rent maximizing city be different from an equilibrium city?³

In an 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

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

³If you've studied microeconomics, you might recognize the parallel between this question and the one answered by the first fundamental theorem of welfare economics. This theorem states that if a market equilibrium exists, then it is Pareto optimal. We will find something similar here.

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 choose \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 blue, pink and grey 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 gray and blue 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 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. This research design provides a basis for valuing location specific amenities that vary discretely as we move

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

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.

Figure 1.10: School district boundaries in Melrose MA around 1990



FIGURE I
Example of Data Collection for One City: Melrose
Streets, and Attendance District Boundaries

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

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), she 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. She finds the A_1 that we want. It ranges between 0.013 and 0.031, so a 1 point increase in test score increase 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 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 were more likely to be college educated and white. Like Black, they found that real estate prices were higher on the high score side of a boundary, but the fact that people were 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 college educated white people, not access to better schools.

1.5 Application #2: Detecting racist property tax assessments

Consider the problem of unfair property tax assessments currently under debate in Chicago. The June 17 edition of the *Chicago Tribune* reports that

An unprecedented analysis by the Tribune reveals that for years

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].*

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 tend to be 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 perfectly plausible that there has been misbehavior by the tax assessor 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 problem is

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

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.

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 blue 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 green 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.

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 blue 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 green line describes the “contract rent

gradient”. This is what the household pays the landlord. This is just enough below the blue 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 very 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 practice, 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%. Each homeowner’s tax bill is the product of the mill rate and their

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

assessed value.

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 out 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.

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?

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

$$\sum_{t=1}^{\infty} \delta^t R^*(x) = \sum_{t=1}^{\infty} \tau \delta^t R_C + \sum_{t=1}^{\infty} \delta^t R_C(x), \quad (1.22)$$

or

$$V(x) = (1 + \tau)V_C(x). \quad (1.23)$$

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

$$\ln V(x) = \ln(1 + \tau) + \ln V_C(x). \quad (1.24)$$

$$\approx \tau + \ln V_C(x). \quad (1.25)$$

Rearranging, we get

$$\tau \approx \ln V(x) - \ln V_C(x). \quad (1.26)$$

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 were to conduct a regression of the tax rate on observed asset prices. What would this look like? Letting i index transaction, it would be something like this,

$$\tau_i = A_0 + A_1 \ln V_C(x)_i + \varepsilon_i. \quad (1.27)$$

and we would 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. That is, while the current system of property tax assessment in Chicago, or in the US as a whole may well be terribly corrupt and unfair. 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, central employment. The open city model makes the following predictions.

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

Second, as commuting costs, t , decrease, utility and consumption, \bar{u} , c^* stay constant (by assumption). As this happens, the rent gradient gets flatter and its intercept stays the same. The extent of the city increases and population increases. We will see some evidence about this later.

Third, as wages, w , increase, utility and consumption, \bar{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 we are 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) 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.

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

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

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,

$$V = \sum_{t=1}^{\infty} \delta^t R. \quad (1.19)$$

Multiplying both sides by δ we get,

$$\delta V = \delta \sum_{t=1}^{\infty} \delta^t R. \quad (1.20)$$

Subtracting equation (1.20) from (1.19),

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

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 $w=3$, $\bar{l} = 1$, $p_c = 1$, $\bar{R} = 0.5$, $\bar{u} = 0$, 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 \bar{x} in terms of w, c^*, \bar{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 \bar{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 \bar{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 $R^*(0)$, \bar{R} and \bar{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.