

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Housing and the Financial Crisis

Volume Author/Editor: Edward L. Glaeser and Todd Sinai, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 978-0-226-03058-6

Volume URL: <http://www.nber.org/books/glae11-1>

Conference Date: November 17-18, 2011

Publication Date: August 2013

Chapter Title: A Spatial Look at Housing Boom and Bust Cycles

Chapter Author(s): David Genesove, Lu Han

Chapter URL: <http://www.nber.org/chapters/c12621>

Chapter pages in book: (p. 105 - 141)

A Spatial Look at Housing Boom and Bust Cycles

David Genesove and Lu Han

3.1 Introduction

This chapter adopts a spatial perspective on supply conditions in housing markets, both within and across urban areas. We investigate empirically how the behavior of prices over the boom-bust cycle of housing markets differs according to the location within the urban area and across urban areas according to the area's size. We take commuting minutes as the appropriate measure of distance from the urban center.

Within markets, we distinguish between households with a longer and shorter commuting time to work. On the presumption that building is easier the further out one gets from the urban center, we expect that a given common increase in demand throughout an urban area will lead to a relatively smaller price response and relative greater quantity response the further away from the center one gets. To test for this, we ask whether households with a shorter commuting trip report relatively higher values for their homes in years when prices are high.

Using the 2005 to 2010 microdata of the American Community Survey (ACS), we find that to be essentially true: controlling for Metropolitan Statistical Area (MSA) cross-year effects, we find that house values decline faster with commuting time in the years 2005 to 2008 than in the bust years of 2009 and 2010. The flatter price gradient of the last two years is consistent

David Genesove is associate professor of economics at Hebrew University of Jerusalem. Lu Han is the Petro-Canada Junior Professor of Business Economics at the Rotman School of Management, University of Toronto.

We thank our discussants, Denise DiPasquale and Stuart Rosenthal, and other participants at the NBER "Housing Market and the Financial Crisis" conference for helpful comments. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c12621.ack>.

with a common demand bust, with spatial variation in supply conditions. This result holds when we consider various subsamples: recent home movers, household heads that work in the central city, and those that commute by car. The result continues to hold when we control for household predicted income, whose effect on house value we also permit to vary by year so as to capture the differential effect of the post-2008 recession across income groups. It is also robust to time varying housing attributes effects.

We then consider the spatial patterns of the growth in rents over the period. We find that rents also decline faster with commuting time in the boom years, although the change in the rate of decline is about half that of prices. That would be suggestive of there being a bubble component along with a “legitimate” increase then dissipation of current demand, except that the difference between the two rates of decline is insignificant.¹

We also examine whether the change in the number of permits over the cycle differs between central and noncentral city counties within MSAs. The difference in the relationships is extremely noisy and we are unable to find any significant difference.

Across markets, we distinguish between cities with longer and shorter average commuting time. The geometry of expanding areas yields a simple property for the supply elasticity: because the area of a circle, or any slice of it, increases with the square of its radius, while the circumference increases linearly, the additional housing built in response to a given price increase is a smaller percentage of the overall stock, the larger is the city. Thus larger cities should have a less elastic supply curve, so that in response to a common demand shock across cities, prices should rise proportionately more and construction less in the larger cities. We argue that the measure of a city should be its mean commuting time.

To check this cross-market hypothesis, we regress the first difference in the yearly FHFA (Federal Housing Finance Administration, formerly OFHEO) MSA-level price indices on year dummies and the interaction of those dummies with (demeaned) average Census 2000 commuting time in the MSA, over the 1975 to 2009 period, using a sample of 258 MSAs. We find a strongly significant positive correlation between the overall housing price change and the slope of the relationship between housing price changes and average commuting time. In other words, when price growth rises on average nationally, it rises relatively more in cities with higher average commuting time. When overall price growth (i.e., in the mean average commuting MSA) increases by 1 percent, price growth in an MSA with a 1 standard deviation greater average commute (2.8 minutes) is predicted to grow one-tenth of a percent more.

1. Like Glaeser, Gyourko, and Saiz (2008), although we use the language of bubbles, nothing we present here can be considered a strict test of the presence of a bubble. Based on the analysis of prices in this chapter alone, a demand story could be told. However, the rapid increase and fall of home prices in the United States around the 2006 peak, with no accompanying movement in general economic conditions until 2008, is certainly suggestive of the presence of a bubble.

We also investigate how the relationship between permits and price growth depends on the average commuting time. We find that log MSA level permits are increasing in price growth, but that the extent of the increase is decreasing in average commuting time: given a 1 percent greater price growth, a 1 standard deviation greater average commute decreases the predicted increase in log permits by 0.3 log points. That effect is more than doubled if we distinguish between periods of nominal price increases and decreases. Together, the quantity and price responses are consistent with MSAs responding differentially to a common demand shock, with supply elasticity declining in the size of the urban area.

We begin with the within-market analysis in section 3.2. We first present our argument in detail, then present the data, the estimation procedure, and then the results. In section 3.3, we follow the same sequence of topics for the cross-market analysis. We relate our analysis and our findings to the existing literature in section 3.4. We conclude in section 3.5.

3.2 Within-Market Spatial Variation in House Price, Rents, and Building Permits

At any given distance from the center of an urban area, we see a mixture of developed and undeveloped land, due to differential expectations across landowners, historical accidents, local zoning variation, topographical variation, and a multitude of other factors. Overall, however, since the price of housing falls with distance from the center, there is likely to be a greater and greater share of undeveloped land the further out one gets. That, in turn, implies a greater possibility of a significant supply response the more distant one is from the center. That this is so empirically is the obvious inference from Burchfeld et al. (2006), who show that building on previously undeveloped parcels of land in and around urban areas is far more common the less developed is the area surrounding that parcel.² Furthermore, where housing prices are lower, building will be less dense, and so there will be a greater opportunity to substitute higher density with lower density. Positive demand shocks that are common across an urban area are therefore likely to be translated mostly into price increases in the city center, with little quantity increase, and translated into more development and little price change near the edge.

There are forces that may act in the opposite direction. Redevelopment closer in to the urban center might be at higher density than previously (the

2. The maps in that paper show new land development occurring overwhelmingly along the edges of the previously developed area. Their figure VI shows an inverted U relationship between the probability of development of a previously undeveloped parcel as a function of the percentage of the surrounding square kilometer that is previously undeveloped. It is the rising part of the relationship that is most relevant for urban areas, as that is where they are located, on average (table X in the same paper).

Burchfeld et al. [2006] data only reveal whether there is any development, not its density). Land use regulations may differ within an urban area, and might be systematically more stringent in the suburbs, especially if higher income people live there and they have a greater willingness to pay for a low density environment; indeed, Glaeser and Ward (2009) have shown that in the greater Boston area, historically less dense areas adopt tougher regulations. Historical development patterns may also mean that inner city structures are more depreciated and so more likely to be replaced; Brueckner and Rosenthal (2009) provide evidence for this. Overall, however, it does appear that supply opportunities are greater further away from the core: Brueckner and Rosenthal (2009) also show that new building is a greater proportion of the existing stock the further out one gets from the urban center.³ In any case, whether the supply response through undeveloped land is sufficiently important to outweigh these other factors is ultimately an empirical question.

Note also that the argument for a greater quantity response and smaller price increase at greater distances from the center cannot be supported by a model of homogenous households. In that case, the required indifference of occupants among all locations implies that preferences alone determine relative prices. Rather, there must be some element of heterogeneity and consequent clustering that permits local supply conditions to matter for relative prices.

Finally, we note that commuting time is the appropriate measure of distance from the urban core if most of the disutility from residing at a distance from the central business district arises from the time resources devoted to going to and from work (White 1988). Housing prices will then decline with distance from the core at a rate given by the decline in commuting *time* from the core, and the share of undeveloped land and supply possibilities should vary accordingly.

3.2.1 Data

The main data source used for the within-market analysis is the public use microdata samples of the ACS for 2005 through 2010. Each year's sample includes roughly 1.3 million owner-occupied and rental housing units. We do not use earlier years' data because in those years geographical identifiers below the state level were not available and the geographical coverage was also much smaller. The variables we take from the survey are subjects' reports on the home value if the unit is owned or annual rent if it is rented, commuting time to work and commuting modes, physical home attributes, utility costs associated with their residence, taxes paid (for owner-occupied

3. Using the ACS data described later to investigate the same relationships uncovered by Brueckner and Rosenthal (2009), but in terms of commuting minutes rather than miles, we have found that building age is decreasing in commuting time, up to a 100 minute commute, and that post-2000 dwellings as a percentage of all dwellings is increasing the commute up to 30 minutes and then more or less constant thereafter.

homes), and the PUMA (Public Use Microdata Area) and state identifiers for both the respondent's place of residence and place of work.⁴

We take the MSA to define the urban area. Although the ACS does not include an MSA identifier, in almost all cases a PUMA can be matched to a unique MSA. We drop observations whose residents reside in a PUMA that does not belong to an MSA. Those with a residence in one of the ten PUMAs that overlap more than one MSA are dropped.⁵ In part of our analysis, we distinguish between PUMAs that contain a central city of the MSA and those that do not. Not all PUMAs have a central city status identifier, however.

For part of our analysis, we restrict the sample to homeowners who report having moved to the given housing unit in the last twelve months. These are about 6.5 percent of all households. The logic of concentrating on such households is that their valuations may be closer to the market value, given that they have just recently purchased their home. However, twelve months is not necessarily "recent" in such a volatile environment.

Our regressor of interest is commuting time, which measures the time it takes to commute from the house to work, as reported by the head of the household. This almost always equals an alternative measure calculated as the difference between the time the respondent reports leaving the home and the time he or she gets to work. Two factors, however, are likely to make reported commuting time a noisy measure of the housing unit's distance in commuting time to the core.

First, households differ in how they commute, and some modes are more pleasant than others. Consequently, prices will drop more quickly with commuting time along a ray along which commuting to the center requires several buses rather than a comfortable subway ride. We handle this in part by including dummy variables for the type used. For a more thorough treatment, we redo our analysis on the subsample for which the household head drives (nearly 90 percent of the total sample).

Second, the household may commute in the direction of the core, but not all the way, or may not commute in the direction of the core, but rather cross-commute; that is, around the circle and not into it. Indeed, Anas and Rhee (2007) show that most commuting is not to the central core. The resulting noise is likely to bias the coefficient on commuting time downwards in magnitude by the reliability ratio; consequently, the difference in measured coefficients across years will also be a downward biased estimate of the true

4. We drop vacant, occupied but neither owned nor rented, mobile home, trailer, boat, RV, and van units, homes without complete kitchen facilities or telephones, Puerto Rico, and units with zero or more than nine bedrooms, or zero rooms. We also drop observations that do not report commuting time (about 2.5 percent of the data) and owner-occupied households that do not report taxes (a little over 1 percent of the data).

5. On average, there are about six PUMAs per MSA. Some large MSAs, such as Los Angeles-Long Beach, contain more than sixty PUMAs.

Table 3.1 Summary statistics of commuting time to work (in minutes)

	Mean	S.D.	25th	Median	75th	<i>N</i>
<i>Owner-occupant sample</i>						
All	30.89	23.83	15	25	40	1,817,977
Recent movers	31.39	23.92	15	25	40	119,031
Drivers	29.82	22.31	15	25	40	1,717,938
<i>Renter sample</i>						
All	28.15	23.00	15	20	35	695,557
Recent movers	26.75	22.07	15	20	30	251,420
Drivers	25.60	20.28	15	20	30	559,617

Source: ACS (2005–2010)

Note: S.D. = standard deviation.

difference. To mitigate such bias, we include a dummy variable for whether the household commutes to the PUMA that contains the central core. More importantly, we redo our analysis for the subsample of households whose head commutes to that PUMA.

Note that although we speak as if there is a single center in the MSA, the presence of multiple centers in a given MSA would not in and of itself affect the quality of the commuting variable as a proxy for supply conditions. With multiple centers, prices will decline at the same rate, in commuting time, from the closest center (in a world of homogenous workers and jobs). The commuting time of households who work in the closest center to their residence will still indicate commuting time distance and so supply conditions. It is the commuting time of workers who work outside of the centers, or at a center other than the one closest to their residence, that leads to the noise in the variable.

The top panel of table 3.1 presents the summary statistics on commuting time for homeowners. The average commuting time is at 30.89 minutes, just a touch over half an hour. The median is 25 minutes. There is a large variation in commuting time across households, as seen in the standard deviation of 24 minutes, and by the 25 minute difference between the twenty-fifth and seventy-fifth percentiles. However, commuting time does not vary much across the samples: the average for the recent mover sample is only half a minute more, and for the driver sample, a minute less, and the twenty-fifth, fiftieth, and seventy-fifth percentiles remain constant across the samples.

Table 3.2's two top panels present summary statistics for the log house value sample by year, with the top panel corresponding to the overall homeowner sample, and the second panel to the recent mover sample. The first column shows the annual difference in the yearly, MSA-wide average log value, averaged across the MSAs in our sample. The second column shows the same as a weighted (by the number of observations in the second year in each MSA) mean across MSAs; thus it is essentially the mean annual differ-

Table 3.2 Summary statistics of house price/rent growth (in %)

	Mean (1)	Weighted mean (2)	Cross-MSA S.D. (3)
<i>All owner-occupants (N = 1,817,283)</i>			
2005–2006	8.85	9.09	6.12
2006–2007	4.56	4.15	4.64
2007–2008	-2.43	-3.21	7.55
2008–2009	-5.30	-7.01	9.32
2009–2010	-2.80	-3.94	6.31
<i>Recent mover owner-occupants (N = 119,023)</i>			
2005–2006	8.39	9.03	14.07
2006–2007	2.98	2.56	15.17
2007–2008	-5.97	-6.04	17.29
2008–2009	-6.68	-9.98	22.94
2009–2010	-6.02	-6.99	19.09
<i>All renters (N = 695,557)</i>			
2005–2006	3.54	3.48	9.07
2006–2007	3.39	4.29	5.03
2007–2008	3.76	3.60	4.88
2008–2009	3.18	3.04	5.20
2009–2010	0.74	0.49	5.43

Source: ACS (2005–2010).

Notes: Column (1) presents the simple mean of the cross-MSA mean of house price growth; column (2) presents the weighted mean, where weights equal the number of observations within each MSA for a given year; column (3) presents the cross-MSA standard deviation around the simple mean.

ence in the average log value in the entire sample. The all owner-occupants panel shows values increasing dramatically between 2005 and 2006 at about 9 percent, increasing further at half that rate over the next year to a peak in 2007, then declining over the next three years by a cumulative 9.5 percent or 14 percent. Half of the decline takes place between 2008 and 2009 alone. The third column shows that there is substantial variation in growth rates across MSAs.

The pattern for recent movers (about 6.5 percent of owner-occupants) is qualitatively similar, with values growing over the first two years and declining over the next three. Quantitatively, however, it differs substantially: the 2006 to 2007 increase is much weaker, and the subsequent declines much greater—the latter by about 3 percent in each year. Here, too, 2008 and 2009 registers the greatest decline, but it stands out less. The standard deviation across MSAs is two or three times as great as that for the overall sample.

What explains the difference between the two panels? It is not differences in geographical composition across MSAs between the overall and the recent mover samples, as columns (1) and (3) show the simple mean and standard deviation across MSAs.

One possible explanation is that in a downturn, recent movers are a self-selected sample of individuals that adjust to the fall in prices. The sample of recent movers is likely to have a disproportionate share of individuals who are insusceptible to factors that hinder selling when prices fall. Their immunity may be due to a lack of loss aversion, or being new buyers (entirely, or to the specific MSA market) who are not “tied” to the recent high prices either as a reference point or through portfolio composition, their having purchased their previous home with little debt. The loss averse may be as unwilling to admit to the loss in value when reporting it as they are unwilling to realize it financially; equity lock-in should not affect a respondent’s assessment of market conditions, but will affect his own reservation price for the property, and that may be what he reports. This would explain why the reported values for the high growth period of 2005 to 2006 are so similar in the two panels, while in the subsequent years, in which there is a slowdown and then a decline in prices, the growth rates are smaller for the recent mover panel.

An alternative explanation is that recent movers are more aware of market conditions than nonmovers. They thus report values that are more up to date and more market-specific than those reported by nonmovers, whose information lags that of market conditions. Nonmovers may also not carefully distinguish between news that reports on their market and on the national market, and so may smooth values across MSAs as well. According to this explanation, nonmovers underestimate the slowdown in the market between 2006 and 2007, do not recognize the extent of the decline over the next three years, and miss the extent of differences in growth rates across the areas.

If we are to use self-reported home values to explore the movement of prices, then we need some evidence that values track prices. Others have explored this issue before (e.g., DiPasquale and Somerville 1995; Kiel and Zabel 1999; and Bucks and Pence 2008), but given the dramatic movements in our sample, and in light of previous evidence that sellers are reluctant to recognize losses, we find it appropriate to reexamine the issue with our sample. We are unable to check the relationship at the sub-MSA level, which is the focus of this part of the chapter, as we lack any such indices, but we can at least do so at the MSA level. We begin with a comparison of the figures in table 3.2 to the annual growth rates of the two leading sets of repeat sale price indices, that of the FHFA and Fiserv Case-Shiller, which we show in table 3.3. All columns show the annual difference in the yearly average of the log of the quarterly or monthly (as appropriate) indices. Averaging over the year is appropriate since the ACS is conducted uniformly over the calendar year. The log transformation recovers the coefficient on the difference between the transaction year dummy and the previous year of sale dummy in the repeat sales regressions on which both sets of indices are based, thus making them comparable to our average of log house values.

Column (1) shows the annual growth rates of the purchase-only, national

Table 3.3 Comparison mean house price growth (in %)

	FHFA all purchase (1)	FHFA all transactions (2)	Unweighted FHFA all- transactions (restricted to ACS sample) (3)	Weighted FHFA all- transactions (restricted to ACS sample) (4)	Fiserv Case-Shiller National (5)	Fiserv Case-Shiller 10-City Composite (6)	Fiserv Case-Shiller 20-City Composite (7)	Unweighted ACS house price growth (restricted to Fiserv Case-Shiller sample) (8)	Weighted ACS house price growth (restricted to Fiserv Case-Shiller sample) (9)
2005–2006	5.89	7.49	7.02	7.51	5.07	7.11	7.30	10.05	9.77
2006–2007	0.22	1.56	1.74	1.45	-4.67	-4.53	-3.92	2.51	2.36
2007–2008	-7.67	-2.25	-3.89	-5.16	-17.20	-18.30	-17.16	-6.50	-5.85
2008–2009	-5.29	-3.56	-4.33	-5.61	-12.16	-13.86	-14.26	-13.47	-12.20
2009–2010	-3.00	-2.33	-3.60	-3.82	0.20	2.08	1.23	-7.36	-7.02

level FHFA index, and column (2) that for all transactions (i.e., including refinancing). In column (3) we have taken the average growth rate of the FHFA MSA indices for MSAs that appear in our sample; in column (4) they are weighted by the incidence of the MSA in our sample. Columns (5), (6), and (7) show the Fiserv Case-Shiller National ten City and twenty City Composite indices. Finally, in Columns (8) and (9) we reestimate the mean and weighted mean from the first two columns in the top panel of table 3.2, using observations with a residence PUMA in one of the twenty Fiserv Case-Shiller cities.

There are a number of observations to be made here. First, the Fiserv Case-Shiller indices show a decline already in 2006 and 2007, while the FHFA indices show a small increase or a leveling off (recall that our data show a moderate increase). In the next two years, the Fiserv Case-Shiller indices show huge declines; the FHFA indices also show declines but they are about 10 percentage points smaller than those of Fiserv Case-Shiller. Over the last pair of years, the FHFA indices continue to decline at a moderate rate, while the Fiserv Case-Shiller indices level off or increase somewhat. All this is well known. What is new is that households report values whose average growth rates are much more like the FHFA than Fiserv Case-Shiller figures. The higher declines in the post-bust period, and the smaller increase in 2006 and 2007 for the all-purchase index compared to the all-transaction index, are also reminiscent of the difference between the recent movers and the overall sample. However, by restricting the ACS sample to the twenty Fiserv Case-Shiller cities (Columns [8] and [9]), we can generate that index's dramatic decline for 2008 and 2009,⁶ although the fit for the other years is not as good.

In table 3.4, we present regressions of the average log house value reported in the ACS for a given MSA and year on the log FHFA transaction index for that MSA and year, and MSA fixed effects. The regressions are weighted by the number of ACS observations in that MSA and year. There are four pairs of regressions, where the second of any pair also includes the lagged index. The four pairs cover the samples we will look: all owners, recent movers, owners who work in PUMAs whose central city status is defined, and owners who commute to a central city PUMA. In the ideal case, we would have a coefficient of one on the current index in the bivariate regression. Except for the recent mover sample, we do in fact obtain coefficients very close to one: 0.96 for the overall sample, and 1.01 for the last two samples. The coefficient of 1.21 for the recent movers sample indicates that those reports substantially "overreact" to the FHFA index. However, recall that the FHFA MSA index includes refinancing transactions, and it is not inconceivable that the assessments that underlie them smooth out purchase prices.

6. That the Fiserv Case-Shiller national index is so similar to its city composites throws some doubt on the representativeness of the underlying sample, the county composition of which is not known. This is especially so since the FHFA has shown that its index can mimic the twenty city Fiserv Case-Shiller composite in part by restricting the sample to those twenty cities.

Table 3.4 Mean house value regressions on FHFA price index

	Overall home owner		Recent mover		PUMAs w/ cent. city def.		Work in central city	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
lnHPI(<i>t</i>)	.96 (.02)	.75 (.03)	1.21 (.05)	1.18 (.06)	1.01 (.04)	.72 (.06)	1.01 (.03)	.79 (.04)
lnHPI(<i>t</i> - 1)		.33 (.02)		.03 (.05)		.38 (.03)		.32 (.04)
RMSE	.040	.028	.079	.079	.049	.041	.082	.077
No. of obs.	1,200	1,200	1,200	1,200	1,199	1,199	1,199	1,199

Note: RMSE = root mean square error.

In all four cases, the lagged price index has a positive and significant coefficient when it is included. For other than the recent mover sample, its share of the sum of the coefficients ranges between .29 to .35; for the recent movers sample, however, it is only .025 of the sum. This reinforces our earlier conclusion that recent movers' valuation are more in line with contemporary prices (or assessments) than that of other households. The sample that looks best when considering the bivariate regression is that of households who work in a central city PUMA, for which the estimated coefficient is insignificantly different from one; but, again, the lagged index gets a large weight when it is included.

Finally, our annual building permit issuance data come from the Census Bureau, which publishes the data for about 18,000 permit-issuing places from 1990 to 2009. We aggregate these place-level data to create the county-level data for the within-market analysis.

3.2.2 Estimation

To compare the within-market house price variation over the years in our sample, we estimate the following relationship:

$$(1) \quad \ln P_{jit} = \sum_{t=2005}^{2010} (I_t \times \text{COM}_j) \gamma_{1t} + \sum_{t=2005}^{2010} (I_t \times \text{COM}_j^2) \gamma_{2t} + X_{jit} \beta + u_{it} + \varepsilon_{jit}.$$

Variable P_{jit} is self-reported home value for household j in MSA i in year t . The vector X_{jit} is a set of housing attributes, including dummies for three sizes of acreage, dummies for the number of bedrooms, whether the house is detached, the number of rooms and building age, and the log expenditures on electricity, natural gas, and water, as well as dummies for ten commuting methods. We also include a full set of MSA \times year fixed effects (u_{it}), as is appropriate for a within-market analysis. These fixed effects control for the overall level of prices in the MSA for that year, so, equivalently, we are regressing the deviation of log value from the mean log value for that MSA and year. This, of course, differences out any MSA-level differences in local

amenities, housing density, urban structure, geographical/regulation barriers, and so forth.

The variable of interest in this regression is COM_j , which indicates the self-reported commuting time. We interact this variable and its quadratic with the year dummies. Our goal is to compare the decline in values of two physically identical homes at different proximities to the urban center, from the boom to the bust in the cycle of the second half of the 2000s.

We focus on the change in the price gradient and not on its level since we are unable to control for *all* differences in physical housing attributes and local amenities. This is likely to impart a positive bias to the effect of distance on the level of values. The monocentric model and its variants predict that per-household housing quantity or quality increases in distance from the center, in response to the decreasing price of a standardized unit. Thus the measured effect of commuting time effect on home values will equal the sum of the effect on the value for some standard unit plus the effect on housing quality/quantity. However, for a given house, if the quality/quantity is relatively constant over time compared to large changes in the per standard unit value, the differences in the measured price gradient over time should come close to measuring the true changes in the price gradient. More formally, let $p_t(x)$ be the log price at distance x in year t , and $h_t(x)$ is log housing quantity/quality at distance x in year t , and specify those relationships as $p_t(x) = \alpha_t + x\gamma_t + u$ and as $h_t(x) = \alpha_t^* + x\gamma_t^* + \eta$. We run the regression

$$(2) \quad v_{it} = (\alpha_t + \alpha_t^*) + x_{it}(\gamma_t + \gamma_t^*) + (u_{it} + \eta_{it}),$$

where $v(x)$ is the log of home value at distance x , that is, $v(x) = p(x) + h(x)$. Basic results from the monocentric city model predict that $\gamma_t < 0$, $\gamma_t^* > 0$; the sign of $\gamma_t + \gamma_t^*$ depends on the utility function. We assume that γ^* is constant over the five years of our sample. Given that assumption, the differences in the measured $\gamma_t + \gamma_t^*$ across years, which we estimate by the difference in the coefficients on the interaction of commuting time and year dummies, capture differences in γ over time. Thus our approach is analogous to difference-in-difference estimation.

Although the long durability of housing—coupled with the ACS's random sampling—makes a constant γ^* a reasonable assumption on which to base the empirical analysis, two factors are a cause of concern. The first is renovation and new construction at a standard, or size, different than the existing housing. If these differ systematically across locations with differing commuting times, our results will be biased. A second concern is sampling noise, especially in the smaller subsamples we use, and especially given that we need to control for MSA year effects. To mitigate the bias, and improve precision, we include the physical home attributes in our regressions. To test for robustness, we also interact them with year dummies.⁷

7. We also considered including log property taxes. In MSAs with market-based assessments, taxes are likely to do an excellent job of capturing variation in value across properties in a mar-

An underlying assumption behind our identification strategy is that the year-to-year growth in house prices is mostly driven by demand shocks, regardless of whether they reflect changes in economic fundamentals or a bubble, and not changes in supply. This assumption, although strong, is consistent with the approach taken in a number of papers, including Glaeser, Gyourko, and Saiz (2008), as well as the general presumption that the evolution of prices in this period constituted a bubble that then burst. Also, although supply-side stories have been offered for the increase in prices in the first part of the 2000s, we are not aware of any that have been offered for the bust.

Finally, although the motivating argument has assumed a common demand shock across the urban area, with a bubble and its bursting the leading explanation of the general pattern of prices over the period of our sample, the post-2007 economic downturn is likely to have affected poorer areas more harshly. Since high income groups live further out, any change in the price gradient over that time period may conceivably be reflecting variation in the deterioration of the labor market, and thus in the demand for housing, across income levels. To account for this, we include predicted income (based on the household's age and occupation group) interacted with the year time dummies. We use predicted and not actual income to avoid including a feedback income effect, whereby a decrease in an individual's housing wealth leads it to work harder.

3.2.3 Results

Baseline Results

Table 3.5 presents the regression described in equation (1). Column (1) presents our baseline regression, using the full sample of owner-occupants. For clarity of presentation, we group pairs of years together in the interaction terms. For example, $COM*0506$ is the product of a dummy variable equal to one if the year is either 2005 or 2006, zero otherwise, and the commuting time (measured in units of 10 minutes). The results predict that in 2005 and 2006, a home with a commuting time of the median 25 will have a value that is $100 \times [-.028 \times 2.5 + .0012 \times (2.5)^2] = -6.4$ percent that of a similar home with a zero minute commute. As noted in the previous section,

ket at a given time, thus soaking up much of the regression error and making the estimates more precise; variations in the tax rate will be captured by the MSA-year fixed effect. Unfortunately, taxes may do too good a job if they also capture variation across time within a market. This will occur if assessments are updated frequently, and are in line with market developments at the sub-MSA level. In the extreme case, if assessed values track changes in the value of individual properties (or more exactly, on average with commuting time), the estimated coefficients on commuting time will not reflect changes in the price gradient, once taxes are included. Taxes will then be masking the very change we are looking for. Had we taxes for a given year, say 2005, this problem would not arise, but the reported taxes are for the year of the survey. If we do include taxes, the estimated rotation of the price gradient falls, but with minimal gain in precision, so that there is no added precision that might compensate for the bias.

Table 3.5 House value regressions (all owner-occupants and recent movers)

	(1)	(2)	(3)	(4)
COM*0506	-0.027 (0.002)	-0.036 (0.0019)	-0.030 (0.004)	-0.033 (0.003)
COM*0708	-0.029 (0.002)	-0.038 (0.002)	-0.030 (0.005)	-0.035 (0.004)
COM*0910	-0.012 (0.003)	-0.026 (0.003)	-0.015 (0.005)	-0.021 (0.005)
COM ² *0506	0.0012 (1.31e-04)	0.0017 (1.18e-04)	0.0014 (2.09e-04)	0.0015 (2.00e-04)
COM ² *0708	0.0013 (1.29e-04)	0.0017 (1.15e-04)	0.0014 (2.63e-04)	0.0016 (2.57e-04)
COM ² *0910	4.34e-04 (1.72e-04)	0.0011 (1.53e-04)	6.81e-04 (3.49e-04)	9.15e-04 (3.39e-04)
Vector of predicted Income*year dummies	No	Yes	No	Yes
Price difference at median 2009-2010	-0.028 (0.007)	-0.058 (0.006)	-0.034 (0.012)	-0.047 (0.011)
Price Difference at Median 2005-2006	-0.059 (0.004)	-0.079 (0.004)	-0.065 (0.008)	-0.072 (0.007)
Difference between 2009-2010 and 2005-2006	0.031 (0.005)	0.021 (0.005)	0.031 (0.010)	0.025 (0.010)
F-stat. (p-value)		74.56 (0.000)		28.19 (0.000)
Sample	Full	Full	Recent	Recent
Observations	1,513,018	1,513,018	104,116	104,116

Source: The home owner sample from the ACS (2005-2010).

Notes: The dependent variable is the log of house value. All regressions include MSA \times year fixed effects and house characteristics. Standard errors in parentheses are clustered at the MSA level. Commuting time is measured in a unit of 10 minutes. *F*-stat. reports the statistics for the joint test that coefficients interaction of predicted income with year dummy equal to zero.

that estimate is likely to be biased. But we are interested in its difference over time. The results predict that in 2009 and 2010, a home with the median commuting time of 25 minutes will have a value that is $100 \times [-0.012 \times 2.5 + .0004 \times (2.5)^2] = -2.8$ percent that of a similar home with a zero minute commute. We take the difference between the 2009 and 2010 difference in values at the median and zero commuting minutes and the 2005 and 2006 difference as our measure of the rotation of the price gradient. As we see in the bottom panel of the table, it is 3.1 percent (rounding is responsible for the discrepancy with our calculation in the text) with a standard error of a mere 0.5 percent.

In column (2) we add predicted income and its interaction with the year pair dummies. Including these variables reduces the estimated rotations by a third, to .021, but it remains significant, with a *t*-statistic of about 4. Closer inspection reveals that the fall comes from including predicted income and

not the interaction with the year pairs (although the interactions are jointly highly significant).

In columns (3) and (4), we restrict the sample to recent movers. As noted earlier, we presume that their assessments will be closer to the market value. In column (3), we obtain very similar results to those in our baseline regression of column (1). The estimates are noisier, but far less than what one would expect from the substantially smaller sample: the standard error on the rotation of the price gradient is only doubled to 1.0 percent. The estimate itself is 3.1 percent, the same that we got for the whole sample. As before, including predicted income and its interaction with the year pairs reduces the estimated rotation, but now only to 2.5 percent, and the difference is such that it would not fail a Hausman test.

Whether or not we restrict the sample to recent movers, all four specifications show that the coefficients on the linear component on commuting time are negative and significant in all years, while the coefficient on the quadratic components are positive. The marginal effect is negative up to at least 105 minutes. Although, as noted earlier, this result is not necessary for our identification strategy, it is nonetheless heartening to see that the estimates are consistent with the underlying assumption that the data are well-represented by a monocentric city model, with much of the variation in housing quantity accounted for.

These regressions constrain the coefficients on the physical home attributes to be constant over time. One possible objection is that the hedonic prices of these attributes might differ across years, and it is these changing prices, along with the correlation of attributes with commuting time, that explains our results. To see if this is so, table 3.6 allows for year-varying coefficients on the lot size and number of bedroom dummies. As one can clearly see, doing so barely changes our core estimates. The largest change is in the rotation of the recent mover sample, which drops from 3.1 percent in table 3.5 to 2.8 percent here. Neither does adding interactions of additional physical home attributes have any noticeable effect on the estimates (regressions not shown).

In order to reduce the noise in the commuting time variable due to cross-commuters and those who commute in the direction of the central city, but only part way, we next restrict the sample to households whose heads work in the central city. This subsample has only 85,200 observations, or about 4.7 percent of the owner-occupants. The subsample is so much smaller because it throws out not only those who do not commute to the central city but also those households whose head works in a PUMA whose central city status is not defined by the Census Bureau.

Table 3.7 presents the results. To establish a comparable baseline, we first consider the sample of observations belonging to a PUMA whose central city status is defined. This is about one-third of the sample. Our baseline regression run on this subsample is presented in column (1). The estimated

Table 3.6 House value regressions (all owner-occupants and recent movers)

	(1)	(2)	(3)	(4)
COM*0506	-0.027 (0.002)	-0.036 (0.0019)	-0.029 (0.004)	-0.033 (0.003)
COM*0708	-0.029 (0.002)	-0.038 (0.002)	-0.030 (0.005)	-0.035 (0.004)
COM*0910	-0.013 (0.003)	-0.026 (0.003)	-0.017 (0.005)	-0.022 (0.005)
COM ² *0506	0.0012 (1.33e-04)	0.0017 (1.19e-04)	0.0014 (2.11e-04)	0.0015 (2.02e-04)
COM ² *0708	0.0013 (1.29e-04)	0.0017 (1.15e-04)	0.0014 (2.62e-04)	0.0016 (2.56e-04)
COM ² *0910	4.65e-04 (1.69e-04)	0.0011 (1.51e-04)	7.56e-04 (3.40e-04)	9.52e-04 (3.34e-04)
Vector of predicted				
Income*year dummies	No	Yes	No	Yes
Price difference at median 2009–2010	-0.029 (0.007)	-0.058 (0.006)	-0.037 (0.011)	-0.049 (0.011)
Price Difference at median 2005–2006	-0.059 (0.005)	-0.078 (0.004)	-0.065 (0.008)	-0.072 (0.007)
Difference between 2009–2010 and 2005–2006	0.030 (0.004)	0.020 (0.004)	0.028 (0.010)	0.023 (0.010)
<i>F</i> -stat.		68.41		24.53
(<i>p</i> -value)		(0.000)		(0.000)
Sample	Full	Full	Recent	Recent
Observations	1,513,018	1,513,018	104,116	104,116

Source: Home owner sample from the ACS (2005–2010).

Notes: The dependent variable is the log of house value. All regressions include MSA \times year fixed effects and house characteristics, and allow the year-varying coefficients on the lot size and the number of bedrooms. Standard errors in parentheses are clustered at the MSA level. Commuting time is measured in a unit of 10 minutes. *F*-stat. reports the statistics for the joint test that coefficients interaction of predicted income with year dummy equal to zero.

rotation effect is 3.8 points, as compared to the 3.1 percent from the corresponding column from table 3.6. Restricting the sample to central city commuters (column [3]) increases the magnitude of all six coefficients on the linear and quadratic commuting times, consistent with the restriction reducing the errors in variable problem. The estimated rotation of the price gradient is roughly doubled, to a significant 7.1 percent, with a *t*-statistic of 5.4. Adding predicted income and its interactions (column [4]) has little effect.

Table 3.8 repeats the same exercise as in table 3.7, except that here we do not pair up years, but show separate effects for each year. This table shows the evolution of the price gradient over time in finer detail. The gradient remains very nearly constant over the first four years of the sample. This is so despite the fact that those four years include periods of growth and decline. Then, in 2009, once the markets are in the depths of the decline, the gradient becomes substantially flatter. The coefficient on the linear term is cut at

Table 3.7 House value regressions (MSAs with defined central city PUMAs)

	(1)	(2)	(3)	(4)
COM*0506	-0.021 (0.003)	-0.029 (0.003)	-0.041 (0.007)	-0.044 (0.007)
COM*0708	-0.021 (0.003)	-0.029 (0.003)	-0.047 (0.008)	-0.049 (0.007)
COM*0910	-0.004 (0.004)	-0.016 (0.004)	-0.009 (0.008)	-0.017 (0.008)
COM ² *0506	0.0010 (1.90e-04)	0.0014 (1.84e-04)	0.002 (5.53e-04)	0.002 (5.08e-04)
COM ² *0708	0.0010 (1.86e-04)	0.0014 (1.75e-04)	0.002 (4.77e-04)	0.002 (4.28e-04)
COM ² *0910	2.42e-04 (2.52e-04)	8.09e-04 (2.29e-04)	7.35e-04 (5.33e-04)	4.84e-04 (5.18e-04)
Vector of predicted Income*year dummies	No	Yes	No	Yes
Price difference at median 2009–2010	-0.009 (0.008)	-0.017 (0.007)	-0.021 (0.017)	-0.039 (0.017)
Price difference at median 2005–2006	-0.047 (0.007)	-0.029 (0.005)	-0.092 (0.017)	-0.097 (0.016)
Difference between 2009–2010 and 2005–2006	0.038 (0.008)	0.012 (0.005)	0.071 (0.013)	0.058 (0.013)
<i>F</i> -stat. (<i>p</i> -value)		46.79 (0.000)		27.42 (0.000)
Sample	Full	Full	Work in central cities	Work in central cities
Observations	452,051	452,051	71,397	71,397

Source: Home owner sample from the ACS (2005–2010).

Notes: Columns (1) and (2) are restricted to observations in MSAs with a central city PUMA. Columns (3) and (4) are restricted to the sample that contains households working in central cities only. The dependent variable is the log of house value. All regressions include MSA × year fixed effects and house characteristics. Standard errors in parentheses are clustered at the MSA level. Commuting time is measured in a unit of 10 minutes. *F*-stat. reports the statistics for the joint test that coefficients interaction of predicted income with year dummy equal to zero.

least in half, and substantially more in the central city commuters’ sample, while the coefficient on the quadratic term drops substantially as well. The year 2010’s gradient looks similar. It is very clear that there is a substantial break in the gradient “series” and that it takes place between 2008 and 2009.

Robustness Checks

An additional concern is that commuting time may not adequately represent the disutility from commuting, which may depend on the means of transportation. We have included dummies for different commuting methods, but that obviously will not control for variations in the marginal disutility. To address this concern more fully, we have run all the previous regressions with the further restriction that the household head commutes

Table 3.8 House value regressions: Separate year interaction effects

	(1)	(2)	(3)	(4)
COM*05	-0.023 (0.005)	-0.030 (0.005)	-0.040 (0.011)	-0.043 (0.011)
COM*06	-0.020 (0.003)	-0.029 (0.003)	-0.044 (0.009)	-0.045 (0.009)
COM*07	-0.022 (0.004)	-0.029 (0.004)	-0.045 (0.008)	-0.047 (0.008)
COM*08	-0.020 (0.004)	-0.028 (0.004)	-0.049 (0.011)	-0.051 (0.011)
COM*09	-0.006 (0.004)	-0.017 (0.004)	-0.005 (0.008)	-0.011 (0.008)
COM*10	-5.59e-05 (0.005)	-0.015 (0.004)	-0.015 (0.010)	-0.026 (0.011)
COM ² *05	0.0011 (0.0003)	0.0015 (0.0003)	0.002 (0.0007)	0.0018 (0.0007)
COM ² *06	0.0010 (0.0002)	0.0014 (0.0002)	0.002 (0.0007)	0.0021 (0.0006)
COM ² *07	0.0012 (0.0002)	0.0015 (0.0002)	0.002 (0.0005)	0.0020 (0.0005)
COM ² *08	0.0009 (0.0002)	0.0013 (0.0002)	0.002 (0.0006)	0.0020 (0.0006)
COM ² *09	0.0005 (0.0003)	0.0009 (0.0003)	-0.0002 (0.0006)	4.73e-05 (0.0005)
COM ² *10	-0.0002 (0.0003)	0.0006 (0.0003)	0.0007 (0.0007)	0.0010 (0.0008)
Vector of predicted Income*year dummies	No	Yes	No	Yes
Price difference at median 2009–2010	-0.007 (0.007)	-0.034 (0.006)	-0.023 (0.013)	-0.042 (0.014)
Price Difference at Median 2005–2006	-0.047 (0.009)	-0.064 (0.006)	-0.093 (0.015)	-0.098 (0.015)
Difference between 2009–2010 and 2005–2006	0.040 (0.009)	0.030 (0.009)	0.069 (0.020)	0.055 (0.021)
<i>F</i> -stat. (<i>p</i> -value)		17.99 (0.000)		7.15 (0.000)
Sample	Full	Full	Work in central cities	Work in central cities
Observations	452,051	452,051	71,397	71,397

Source: Home owner sample from the ACS (2005–2010).

Notes: Columns (1) and (2) are restricted to observations in MSAs with a central city PUMA. Columns (3) and (4) are restricted to the sample that contains households working in central cities only. The dependent variable is the log of house value. All regressions include MSA × year fixed effects and house characteristics. Standard errors in parentheses are clustered at the MSA level. Commuting time is measured in a unit of 10 minutes. *F*-stat. reports the statistics for the joint test that coefficients interaction of predicted income with year dummy equal to zero.

to work by driving. None of our results change in any substantive way, which is to be expected given that nearly 90 percent of all owner-occupants commute by car to work.

In comparing across years, the analysis has thus far implicitly used the nationwide average movement in prices or values as indicators of the state of overall market conditions. However, there are differences across cities in the extent and timing of the price appreciation and depreciation over our period of analysis. We thus check to see whether using the MSA level FHFA house price indices instead of year dummies generates the same results. To do so, we replace the interaction of commuting minutes and year dummies with the interaction of commuting minutes and the growth in the MSA price index.

Table 3.9 shows the resulting regression for the sample of households whose head works in the central city. The first column controls for MSA cross-year fixed effects and physical house attributes, as usual. The coefficient of $-.041$ in column (1) implies that when prices grow at a 1 percent lower rate, the absolute slope of the price gradient decreases by .04 percent per commuting minute. Thus prices at a 25-minute commute fall one-tenth of a percent more than prices at the center for every 1 percent decline in overall price.

At a 16 percent decline, which corresponds to the decline between 2006 and 2010, according to both the valuations reported in the ACS and the FHFA All Purchase Index, the estimated rotation is thus 1.6 percent. That is exactly one-half of our baseline estimate, and less than a quarter of that

Table 3.9 House value regressions on FHFA price indices (work in central city sample)

	(1)	(2)	(3)	(4)
COM	-0.016 (0.002)	-0.015 (0.002)	-0.013 (0.000)	-0.016 (0.002)
COM*($\Delta \ln \text{HPI}$)	-0.041 (0.023)	-0.044 (0.019)		
COM*($\Delta \ln \text{HPI}$)* $I_{\{\Delta \ln \text{HPI} > 0\}}$			-0.075 (0.031)	-0.021 (0.025)
COM*($\Delta \ln \text{HPI}$)* $I_{\{\Delta \ln \text{HPI} < 0\}}$			-0.007 (0.035)	-0.069 (0.028)
$\ln \text{HPI}$		1.094 (0.054)		1.097 (0.056)
Lagged $\ln \text{HPI}$		-0.062 (0.089)		-0.054 (0.089)
MSA \times year fixed effects	Yes	No	Yes	No
MSA fixed effects		Yes		Yes

Source: ACS (2005–2010).

Notes: The dependent variable is the log house value. “HPI” is the MSA FHFA house price indices. Standard errors in parentheses are clustered at the MSA level. Commuting time is measured in units of 10 minutes. The number of observations is 85,143.

from the work in central subsample. Using the Fiserv Case-Shiller decline of 31 percent brings the estimate in line with that of the baseline estimates, but not quite one-half that of the work-in-central-city sample.

In column (2) we replace the MSA cross-year fixed effects with the price index and its lag and MSA fixed effects. The estimates on the commuting terms remain essentially unchanged. In the next two columns we allow the gradient to shift with a differential magnitude according to whether growth is positive or negative, by replacing the interaction of commuting time and MSA price growth with the following two variables: $COM \times \Delta \ln HPI \times I\{\Delta \ln HPI > 0\}$ and $COM \times \Delta \ln HPI \times I\{\Delta \ln HPI \leq 0\}$. Both the durability of housing and our earlier result that the price gradient shifted when prices fell over the latter part of the 2005 to 2010 period but not when they rose suggest that allowing for such an asymmetry might be important. We obtain ambiguous results: when MSA cross-year effects are included (column [3]), we find that the price gradient shifts most when prices are rising; when the price index is included instead, we find the opposite.

Within-Market Spatial Variation in Rents

We now turn to consider the rents of rented dwellings. Unlike house prices, for which expectations over future conditions matter, rents are determined solely by current demand and current supply. The comparative behavior of rents to prices over a boom and bust cycle can thus give us some sense of whether the price movement was the result of a positive shock to current demand for housing services that then dissipated, or a bubble. (Of course, an increase in expected future demand that is then undone is indistinguishable from a bubble.) If it is a temporary shock to current demand, then the rent gradient should become flatter in the bust, just as for prices. If it is a bubble, then the spatial pattern of rents should remain constant over time. That is, unless the boom has persisted long enough for a substantial overhang to have developed, in which case rents should fall more where supply is more elastic. Thus, for a bubble, the rent gradient should either remain constant over time, or become steeper in the bust.⁸

We first turn back briefly to table 3.2, where the bottom panel presents the annual average growth in rents. Here, rents behave very differently than self-reported values and price indices. In contrast to the boom and bust cycle displayed by the latter, rents increase steadily year after year, at about 3.5 percent a year, until 2010, when they essentially stagnate at the 2009 level. This is clearly at odds with any explanation of the boom and bust of prices over the 2000s based on changing current demand for housing services.

Table 3.10 presents results from regressing log rents on the same set of

8. Obviously, any overhang will affect prices as well; however, since prices will incorporate the gradual disappearance of the overhang with the normal growth in demand, the effect on prices should be substantially less.

Table 3.10 Rent regressions (work in central city sample)

	(1)	(2)	(3)	(4)	(5)	(6)
COM*0506	-0.026 (0.009)	-0.029 (0.008)	-0.035 (0.010)	-0.036 (0.010)	-0.002 (0.012)	-0.006 (0.012)
COM*0708	-0.014 (0.008)	-0.017 (0.007)	-0.021 (0.008)	-0.024 (0.007)	-0.008 (0.011)	-0.003 (0.010)
COM*0910	-0.004 (0.009)	-0.013 (0.008)	-0.017 (0.009)	-0.024 (0.008)	0.014 (0.014)	0.004 (0.001)
COM ^{2*} 0506	6.95e-04 (4.66e-04)	0.0016 (5.85e-04)	0.0020 (6.64e-04)	0.002 (6.56e-04)	-6.96e-04 (9.04e-04)	-3.48e-04 (9.46e-04)
COM ^{2*} 0708	3.17e-04 (4.66e-04)	8.38e-04 (4.21e-04)	9.30e-04 (5.17e-04)	0.001 (4.77e-04)	-2.25e-04 (6.33e-04)	9.57e-05 (5.84e-04)
COM ^{2*} 0910	3.17e-05 (6.94e-04)	8.27e-04 (6.03e-04)	0.0011 (5.77e-04)	0.001 (5.77e-04)	-8.64e-04 (0.0011)	-1.88e-04 (9.72e-04)
Predicted income with year interactions	No	Yes	No	Yes	No	Yes
Price difference at median 2009–2010	-0.008 (0.019)	-0.027 (0.016)	-0.036 (0.019)	-0.051 (0.017)	0.030 (0.029)	0.009 (0.023)
Price difference at median 2005–2006	-0.056 (0.018)	-0.061 (0.017)	-0.076 (0.022)	-0.076 (0.021)	-0.021 (0.023)	-0.017 (0.024)
Difference between 2009–2010 and 2005–2006	0.048 (0.019)	0.034 (0.018)	0.040 (0.025)	0.025 (0.026)	0.051 (0.025)	0.026 (0.024)
<i>F</i> -stat.		5.17		0.66		6.04
(<i>p</i> -value)		(0.007)		(0.52)		(0.003)
Sample	All rentals	All rentals	Rented apartments	Rented apartments	Rented houses	Rented houses
Observations	25,182	25,182	14,975	14,975	10,207	10,207

regressors as in tables 3.5, 3.6, and 3.7: year dummies, their interaction with commuting time and with the square of commuting time, and attributes of the unit. The sample is restricted to households whose heads commute to the central city. The first two columns show the results on the sample of all renters, with the interactions of predicted income and year dummies controlled for in column (2) but not so in column (1). We first note that the rent gradients are downward sloping over the vast majority of the support of commuting time. In the early years, commuting time decreases rents until 105 minutes, similarly to what we saw for values; in 2009 and 2010, however, it turns up at 51 minutes (the ninety-fourth percentile), although the quadratic term is insignificant. The estimation rotation is 3.9 percent, and is significant. Thus we can reject the pure bubble hypothesis, under which rents would have fallen equally along the commuting time dimension, or fallen less at the center due to overhang. A weaker test is to check whether rents rotate up with the housing bust less than prices themselves. The point estimate for the rotation is about half that of the 7.1 percent we obtained in table 3.7 for the parallel sample of owner-occupants who commute to the central city, but the difference between the two is insignificant.

The next four columns distinguish between rented apartments (columns [3] and [4]) and rented houses (columns [5] and [6]). The distinction might be an important one, since nominal rigidity in rents is much more prevalent in single-family dwellings than in apartments, likely because apartments are much more likely to be owned by corporations, partnerships, and large investors (Genesove 2003). Thus we might expect that the gradient for houses would be relatively constant over time. Nonetheless, we find a positive rotation of the rent gradient in both categories, with that for houses much bigger and significant only.

Thus, surprisingly, although the aggregate behavior of reported rents is very different from that of both reported values and price indices, the variation in the temporal behavior of rents across commuting time is qualitatively similar to that of prices. Taken by itself, the similarity in the spatial behavior is consistent with a dramatic temporary increase in current demand in the bust period that gets undone between 2008 and 2009. It is not consistent with either a bubble that gets pricked, or an increase in future expected demand that then goes away; in such cases, the rent gradient would have stayed constant or because of the overhang, become steeper.

Within-Market Spatial Variation in Building Permits

Although the focus of this chapter is on prices, since our analysis does presume that building in the city center is more difficult than building outside it, we now turn to ask whether building in the urban center is indeed less responsive to price increases. We aggregate up county-level permit data to the set of central city counties and noncentral city counties, in each MSA, for those counties for which permit data are available. We consider only those

Table 3.11 Summary statistics of permits

	(1)	(2)	(3) = (1) – (2)	(4)	(5)
	Log permits		Central- noncentral difference	First difference of central-Noncentral difference	
	Central	Noncentral		Mean	S.D.
2005	6.36	7.60	1.23	—	—
2006	6.22	7.43	1.21	-0.02	0.30
2007	5.93	7.17	1.24	0.03	0.32
2008	5.48	6.71	1.23	-0.01	0.30
2009	5.17	6.43	1.26	0.03	0.38

Notes: The building permits are obtained at the county level from the Census (1990–2009) and aggregated up to the set of central and noncentral counties in each of the 141 MSAs that have both central and noncentral counties.

141 MSAs for which we are able to obtain permit data for both central and noncentral city counties. In columns (1) and (2) of table 3.11 we present the log of the number of permits in the central city counties and the noncentral city counties, respectively, averaged across the MSA by year, from 2005 to 2009, with the difference given in column (3). Because the coverage of counties is not complete (although the sample is balanced), and the number of permits is not normalized by the existing housing stock, we cannot infer the relative degree of building in these two sets of areas. Column (4) shows the first difference in these figures, however, and thus shows us how the extent of building moves differentially over the cycle. Column (5) shows the standard deviation of these changes across MSAs. With 141 MSAs, the standard error on the mean first differences in column (4) are thus about 0.03. There is clearly no evidence that the relative number of permits within the urban areas varies over the period.

In table 3.12, we ask whether the number of permits increases less with price growth in central city counties than in noncentral city counties, as would be implied by a smaller supply elasticity closer in to the urban center when demand shocks dominate supply shocks. That it is price growth that should be expected to be correlated with permits and not the price levels follows naturally from the fact that permits essentially indicate the change in supply. Mayer and Somerville (2000) show that this logic follows from the standard monocentric city model, specifically Capozza and Helsley (1989). To account for replacement investment, we nonetheless allow for a price level effect by adding the current price level as well and not simply the difference. Underlying this investigation is of course an assumption that supply shifts are substantially less variable than demand shifts.

Column (1) shows the regression of the within-MSA log difference in permits between central and noncentral city counties on the MSA-level price

Table 3.12 Permit regressions

	(1)	(2)	(3)	(4)	(5)	(6)
lnHPI	0.27 (0.30)		-0.16 (0.15)		0.20 (0.20)	
Δ lnHPI	1.32 (1.15)	1.29 (1.11)	0.28 (0.48)	0.22 (0.49)	0.18 (0.51)	0.25 (0.51)
Trend					-0.013 (0.004)	-0.010 (0.003)
Sample period	2005–2009	2005–2009	1990–2009	1990–2009	1990–2009	1990–2009
Number of obs.	585	585	2,222	2,222	2,222	2,222

Notes: The MSA fixed effects are included in each regression. Standard errors are in parentheses.

index and its first difference, over the 2005 to 2010 period. Neither variable is close to being significant. In column (2) we drop the price level, but this does not improve the precision on the first difference. Columns (3) and (4) repeat the estimation over the entire period for which we have data, from 1990 to 2010, and columns (5) and (6) add a trend term. In none of the regressions is there a statistically significant relationship.

3.3 Cross-Market Spatial Variation in House Price and Building Permits

So far, our analysis has focused on the spatial variation in house price growth among housing units within the same market, where the variation is collapsed to the time to commute dimension. It is natural to extend the use of commuting time as a proxy for supply conditions to the cross-market level. In this section, we first derive a theoretical relationship between average commuting time and housing supply elasticity across markets. It is based on the geometry of expanding areas. We then empirically test this relationship by examining how house prices and building permits vary across markets with different average commuting time, and how these effects vary with the stages of the housing market cycle.

The intuition for the theoretical result is shown in figure 3.1. If all the additional housing built in response to a price change is built on the margins of a circular city, the increased stock will equal the circumference of the circle. As the total housing stock in the city is the area of the circle, the percentage change in the housing stock for a given absolute price change should be decreasing in the city radius. In considering the supply elasticity, we need to consider a given percentage increase in price, of course, and prices are higher in bigger cities. But it is only the location rents component of prices that increases with city size, and in Capozza and Helsley's (1989) dynamic version of the monocentric city model, which we base our analysis on, it increases linearly. The presence of the remaining components of price, agricultural rents, and construction costs, ensure then that supply elasticity




Fig. 3.1 Graphical illustration of supply elasticity with respect to city size

is inversely proportional to the city radius. As the following analysis also shows that the proper measurement of the radius is maximum commuting time, not distance per se, we show that supply elasticity is inversely related to commuting time.

In developing our argument, we follow Mayer and Somerville (2000) and Green, Malpezzi, and Mayo (2005) in deriving an elasticity of supply based on Capozza and Helsley (1989). We view our presentation as somewhat more transparent, however. Consider, then, a set of cities that have developed in line with that model. The cities are of different sizes, perhaps because they were established at different points in the past, or perhaps because they have different transportation infrastructures that yield different commuting costs.

Now consider a shock to demand that takes the form of an enhanced willingness to pay in the city. The greater price will induce an increase in the housing stock. In line with Capozza and Helsley (1989), all new development takes place at the city edge. (In practice, redevelopment in the city core is certainly possible. However, as discussed in section 3.2, the empirical literature generally shows that supply opportunities are much greater at the edge.)⁹ Assume that households are willing to pay α percent above the current fundamental price at the city edge. That means an absolute price increase at the city edge of $\Delta P = \alpha P_E$, where P_E is the current price at the edge. Capozza and Helsley (1989) show that

$$P_E = A + C + Tf(g, r)z_E,$$

where A is the discounted value of land use in its undeveloped state (agricultural rent) and C is the conversion cost, which we should see here as the construction cost of the housing structure. The last term represents the

9. To control for the possibility that some cities develop more in the core than others, our empirical specifications include MSA fixed effects and the interactions between year dummies and supply constraint proxies, such as the Wharton Residential Land Use Regulation Index and the undevelopable land share.

discounted value of location rents, with T the time cost of commuting a unit distance, z_E the distance from the center to the edge, g population growth, r the interest rate, and f a function that is increasing in g and decreasing in r .

Define the absolute rate at which prices decline with distance from the city center as k . Then the city edge grows out an additional $\Delta z = k^{-1} \Delta P$ from the center. The additional area that is developed is $2\pi z_E \Delta z = 2\pi k^{-1} \alpha \{ [A + C] z_E + Tf(g, r) z_E^2 \}$. (See figure 3.1. Like the aforementioned papers, we assume a constant lot size.) Since the area of a circular city is z_E^2 , we obtain that the percentage increase in the housing stock is $2k^{-1} \alpha \{ [A + C] z_E + Tf(g, r) \}$. Since $k = T/r$ (a dwelling at one unit of distance closer to the center is worth the discounted value of the time cost of commuting a unit distance more), the supply elasticity of housing is

$$(3) \quad \eta = 2r \left\{ \frac{[A + C]}{(Tz_E)} + f(g, r) \right\}^{10}$$

As is often done, we will assume that the opportunity cost of the land plus the construction cost, $A + C$, do not vary substantially across cities. The housing supply elasticity is thus decreasing in the maximum commuting time (Tz_E), increasing in future population growth g , and increasing in the sum of the opportunity cost of land and construction costs ($A + C$).

Our focus is on commuting time. Importantly, the previous argument shows us that the proper measure of a city's size as a factor of supply elasticity is commuting minutes and not kilometers. The term that appears in the supply elasticity is the maximum commuting time, but given the greater sensitivity of the maximum to measurement error, we substitute the average commuting time, which is two-thirds of the maximum under this model.¹¹

If demand shocks are a national phenomenon (as Cotter, Gabriel, and Roll [2011] show to be increasingly the case) that add a willingness to pay to inverse demand of a constant percentage over current prices, then the foregoing implies that prices should rise relatively more in larger cities. In smaller cities, they will be undone by massive building along the edge, which will keep prices from increasing too much.

3.3.1 Data

To test these implications, we look at the differential movement of MSA-level house prices across time according to the average commuting time of

10. Green, Malpezzi, and Mayo (2005) derive the semi-elasticity of housing supply from the Capozza-Helsley model as $2(r - g)/(Tz_E)$. Multiplying this by P_E yields our result. Assuming $g = 0$, our expression also reduces to that implicit in Saiz (2010). His central claim is that the supply elasticity is decreasing the undevelopable share of the city. Given the population, this is so: to fit in the same number of households in a semicircle as in a circle, the former must extend further out from the origin, and so by precisely the argument given here, must have a smaller supply elasticity.

11. With a fixed lot size, the mean is proportional to the maximum: since the number of dwellings at distance z from the center is $2\pi z$, mean commuting time is $2\pi T \int_0^{z_E} z^2 dz / 2\pi \int_0^{z_E} z dz = 2/3 T z_E$.

the MSA. We measure house prices with MSA-level FHFA all transaction (i.e., purchase plus refinancing) indices. This covers a somewhat different set of MSAs than the ACS data. We also consider a much longer window: 1976 to 2009. The sample of 257 MSAs is heavily unbalanced due to different starting times, which are distributed between 1975 and 1993. While the price series is at a quarterly frequency, we use the averaged annual data, since our focus is on patterns over a cycle, not on higher frequency, seasonal effects. As before, we obtain the housing permit data from the US Bureau of the Census. We aggregate across places to create metropolitan-area-level aggregate permits. Annual permits data are available for 1990 to 2009. In some regressions, we also control for measures of growth to control for g . We also at times condition on the interaction with year dummies of other proxies for determinants of housing supply, notably the updated Wharton Residential Land Use Regulatory Index (WRLURI) (Gyourko, Saiz, and Summers 2008) and Saiz's undevelopable land share (Saiz 2010).

To measure commuting time in each location, we use the average number of minutes needed for a one-way trip to work among workers sixteen years and over from the 2000 Census. The mean average commuting time is 22 minutes, with a standard deviation of 2.66 minutes. The MSAs with the longest average commuting time are New York, NY (37.4 minutes); Newburgh, NY-PA (34.6); Bremerton, WA (34.4); Monmouth Ocean, NJ (33.8); and Washington, DC (33.5). Those with the shortest are Grand Forks, ND-MN (16.0 minutes); Dubuque, IA (16.7); Waterloo-Cedar Falls, IA (17.1); Great Falls, MT (17.2); and Bismarck, ND (17.4). Figure 3.2 reports the histogram of average commuting time. Note that our measure of average commuting time does not vary over time; there should therefore be little concern of an endogeneity bias in which positive demand shocks both increase prices and increase congestion and thus commuting time.

3.3.2 Cross-MSA Variation in House Prices

The brief argument we laid out implies that when demand increases, price will increase more in markets with longer average commuting time. The argument has nothing to say about demand decreases (nor has the underlying Capozza and Helsley [1989] model), but fortunately, most of the period of our sample is one of price increases. Coupled with the assumptions that (a) demand shocks are much more variable than supply shocks, and (b) demand shocks are heavily correlated across MSAs (Cotter, Gabriel, and Roll 2011), the claim implies that when overall prices increase, prices will increase more in high average commuting MSAs.

Our procedure has two stages. In the first stage, we estimate the following regression model:

$$y_{it} = \sum_{t=1975}^{2009} I_t \alpha_t + \sum_{t=1975}^{2009} (I_t \times \text{COM}_j) \beta_t + \sum_{t=1975}^{2009} (I_t \times X_j) \gamma_t + u_i + \varepsilon_{it}.$$

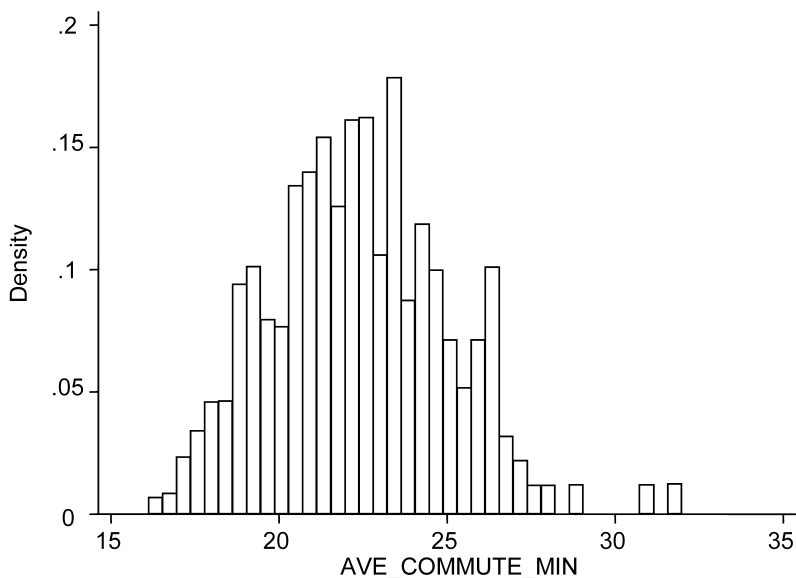


Fig. 3.2 Histogram of average commuting time (mean = 22.39 min., sd = 2.66 min.)

The dependent variable is the first difference in the log of the house price index in MSA i at year t .¹² Variable I_t is a dummy for year t , COM_i indicates the demeaned, average commuting time in MSA i from the year 2000 Census, and X_j indicates other possible determinants of the supply elasticity. Finally, MSA fixed effects are included. These effects control for, among other things, trends in market-specific geographical feature and regulatory constraints, and unchanging amenities.

With demeaning and MSA fixed effects, the parameter α_t reflects the overall housing price growth in year t . The parameter β_t reflects, in a given year t , how the price growth varies across markets with different average commuting time. Figure 3.3 shows the time series of the estimated parameters, that of the raw estimates on the left and that of the residuals from regressions of each on controls to be included in regressions below on the right.

In the second stage, we explore the relationship between β_t and α_t . If longer average commuting time indeed proxies for more inelastic housing supply, then we should expect that when overall house prices increase, house prices will increase more in the MSAs with longer commuting time than in others; that is, that β_t and α_t should be positively correlated over time. We

12. The index is equal to $100\exp(\psi_t)$, where ψ_t is the estimated coefficient on the year t variable that takes values $\{1, -1, 0\}$ according to whether t is the year of sale, previous sale, or neither, in the regression with log price as the dependent variable. Thus taking the log of the index yields ψ_t , the measured price growth, plus a constant.




Fig. 3.3 Cross-MSA house price analysis

Note: Estimated α_t and β_t , from the Regression of First Difference of log Price Index on Year Dummies and Interaction of Year Dummies with Average Commuting Time.

wish to examine this relationship through the regression $E[\beta_t|\alpha_t] = a_0 + a_1\alpha_t$.¹³ The argument we outlined earlier implies that a_1 should be positive.

Of course, we do not observe β_t and α_t , but only their estimates $\hat{\beta}_t$ and $\hat{\alpha}_t$. We thus face a measurement error problem, which implies that the ordinary least square (OLS) estimate of a_1 will be biased: it converges to the sum of the attenuation of the true value (due solely to the measurement error in $\hat{\alpha}_t$) and an additional bias of the sign of the covariance between the estimation errors of the two coefficients. However, since we can estimate the distribution of those errors, using the standard asymptotic results on the distribution of OLS estimates, we can form the following consistent method of moments à la Fuller (1987) and Buonaccorsi (2010):

$$\alpha_1^{MOM} = \frac{\{S_{\alpha\beta} - Q_{\alpha\beta}\}}{\{S_{\alpha\alpha} - Q_{\alpha\alpha}\}},$$

where $S_{\alpha\beta}$ is the sample covariance of $\hat{\beta}_t$ and $\hat{\alpha}_t$, $S_{\alpha\alpha}$ is the sample variance of $\hat{\alpha}_t$, $Q_{\alpha\beta} \equiv (1 + 1999 - 1976)^{-1} \sum_{t=1976}^{1999} \sigma_{\alpha\beta t}^2$, $Q_{\alpha\alpha} \equiv (1 + 1999 - 1976)^{-1} \sum_{t=1976}^{1999} \sigma_{\alpha\alpha t}^2$, $\sigma_{\alpha\alpha t}$ is the estimated standard error on $\hat{\alpha}_t$, and $\sigma_{\alpha\beta t}^2$ is the estimated covariance of $\hat{\alpha}_t - \alpha_t$ and $\hat{\beta}_t - \beta_t$. Where there are no other regressors other than the year dummies and their interaction with commuting time, $\sigma_{\alpha\beta t}^2$ equals minus the regression error variance time the product of the mean of the commuting variable divided by the variance of the same, where the mean and variance are taken over the set of MSAs in the sample for that year. Since we use the deviation of average commuting minutes from its mean, $Q_{\alpha\beta}$ is very nearly zero (it is not exactly zero as the sample is not balanced). Thus in our case

13. We mean this as not only a projection, and attribute no casual interpretation to it. We estimate a regression coefficient since correcting it for errors in variables is more straightforward than correcting a correlation estimate.

the method of moments estimator corrects the OLS estimator essentially for attenuation bias (the error in α).¹⁴

Table 3.13 reports the estimates from the second-stage regressions. Since the first stage uses the growth rate in the nominal housing price index, α , includes the inflation rate, and so we first adjust $\hat{\alpha}$, by subtracting off the inflation rate. In the baseline specifications of column (1), which shows the “naïve” (OLS) coefficient, and column (2), which shows the method of moments estimator, we see that the coefficients on (the adjusted) $\hat{\alpha}$, are positive and statistically significant at the 1 percent level, consistent with our hypothesis. The coefficient of 0.047 implies that in a year in which house prices increase on average by 1 percent, the price in an MSA whose average commuting time is 2 standard-deviations longer than the sample mean (that is, 28 minutes instead of 22.4 minutes) will increase by about 26 percent ($0.047 * (28 - 22.4) \times 100$ percent) more than the average. The method of moments estimator is nearly exactly the same as the naïve estimator, since here the attenuation bias is small, as the year effects are precisely estimated relative to the large changes in overall housing prices over time.

The remaining columns in table 3.13 show that the positive and significant relationship between $\hat{\beta}_1$ and $\hat{\alpha}$, is robust to a number of specification changes. First, in column (3), we include the year variable in the second-stage regression to control for the time trend; the trend is insignificant, and the estimated commuting time effect is unaffected. Second, we address the concern that the first-stage regression for the baseline specification does not adequately control for other conditions that affect the elasticity of supply. In the next three columns, we consider the Wharton land regulation index and Saiz’s undevelopable land share, by adding their interaction with the year dummies to the set of regressors. Column (5) shows the method of moments estimator, for which the estimated regression coefficient of $\hat{\beta}_1$ on $\hat{\alpha}$, is more than one-half that of previously, although it remains significant. Once again, adding a trend has no effect on the OLS estimate. Another determinant of supply elasticity is growth, as equation (3) shows. In columns (7) through (9) we add the interactions of population growth and average income growth to the set of regressors. The method of moments estimator is again reduced relative to the baseline, although less so, and it remains positive and significant. In the last set of columns, we add in both sets of alternative supply constraints. The results are like those in columns (4) through (6), where only the land regulation index and the undevelopable share are included.

Thus we can conclude that when prices rise overall, they rise more in high average commuting cities—a result predicted by the simple geometry of expanding areas, in the context of highly correlated and similarly sized MSA-level demand shocks. This result holds even when controlling for

14. This is not the case when we include other regressors. However, in practice, most of the bias comes from the attenuation component.

Table 3.13 Second-stage of the cross-MSA price variation analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\hat{\beta}_t$												
$\hat{\alpha}_t$	0.047 (0.020)	0.050 (0.004)	0.050 (0.020)	0.012 (0.014)	0.020 (0.003)	0.014 (0.014)	0.025 (0.015)	0.035 (0.003)	0.030 (0.015)	0.003 (0.010)	0.012 (0.002)	0.007 (0.009)
Time trend	No	No	-0.0001 (0.0001)	No	No	-0.00004 (0.00005)	No	No	-0.00009 (0.00006)	No	No	-0.00009 (0.00004)
Population and income controls	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Other supply constraints	No	No	No	Yes	Yes	Yes	No	No	No	Yes	Yes	Yes
Estimation method	OLS	MOM	OLS	OLS	MOM	OLS	OLS	MOM	OLS	OLS	MOM	OLS

Notes: The independent variable is $\hat{\alpha}_t$, the estimated coefficient on $I_{t,r}$, the dummy variable for year t , adjusted for inflation. The dependent variable is $\hat{\beta}_t$, the estimate of the coefficient on the interaction of the (demeaned) average commuting time with $I_{t,r}$. Data sources for first stage: (1) MSA-level FHFA Office of Federal Housing Enterprise Oversight (OFHEO) house price indices (1976–2009), time averaged over the year. (2) Average commuting time from the 2000 Census data. The first-stage regression regresses $\ln HPI$ or $\Delta \ln HPI$ on year MSA fixed effects, dummies, and the interactions of average commuting time and year dummies. “Population and income controls” refers to the interaction of population growth and average income growth with year dummies. “Other supply constraints” refers to the interaction of the WRLURI and Saiz’s undevelopable land share with year dummies. MOM = method of moments.

regulatory and topological supply proxies that have previously been used in the literature.

3.3.3 Cross-MSA Variation in Building Permits

We now ask whether building activity increases more with prices in markets with low average commuting time, as we argued earlier. To test this, we estimate the following regression:

$$\ln\text{PERMIT}_{it} = \alpha_0 \ln\text{HPI}_{it} + \alpha_1 \Delta \ln\text{HPI}_{it} + \beta \Delta \ln\text{HPI}_{it} \times \text{COM}_i + \gamma' X_{it} + e_i + \epsilon_{it},$$

where i indexes the metropolitan areas, t represents the years from 1990 to 2009, COM_i indicates the average commuting time in MSA i , $\ln\text{HPI}_{it}$ indicates the inflation-adjusted log of the FHFA house price index in MSA i in year t , $\Delta \ln\text{HPI}_{it}$ is its first difference, X_{it} is income and population and their growth rates, and e_i is an MSA fixed effect.

The coefficient α_1 indicates how construction activities are associated on average with price changes, while β captures how that association varies with average commuting time. Since, according to the argument we developed earlier, the elasticity of supply is decreasing in average commuting time, we expect β to be negative. As before, we expect permits to be most strongly correlated with the price changes rather than its level (Mayer and Somerville 2000), but we also include the price level to account for depreciation.

We present the estimates in table 3.14 using single-family permits (the estimates for total permits are very similar). Column (1) shows the regression without income and population. It shows a highly significant positive association between the price change and the level of permits, with every 1 percent increase in price growth associated with about a 4 percent increase in permits. As predicted by our theoretical argument, that relationship is

Table 3.14 Permit regressions

	(1)	(2)	(3)	(4)	(5)
	<i>Single-family permits</i>				
$\Delta \ln\text{HPI}$	0.02 (0.12)	0.12 (0.13)	0.13 (0.13)	0.41 (0.14)	0.06 (0.17)
$\Delta \ln\text{HPI}$	4.23 (0.19)	4.21 (0.19)	4.09 (0.18)	1.87 (0.33)	5.89 (0.55)
$\ln\text{HPI} \times \text{COM}$	-0.095 (0.046)	-0.071 (.043)	-0.074 (0.044)	-0.146 (0.077)	-0.25 (0.12)
Other controls		Income, population	Income, population Δ Income Δ Pop	Income, population	Income, population
Observations	2,559	2,559	2,559	1,173	846

Note: The building permits are obtained at the MSA level from the Census (1990–2009). All regressions include MSA fixed effects.

decreasing in the average commuting time in the MSA: a 1 standard deviation (2.8 minute) greater average commute decreases the predicted increase in single-family home permits given a 1 percent greater price growth rate from 4.16 percent to $(4.16 - 2.8 \times 0.095) \times 100$ percent = 3.89 percent. The price level is irrelevant. The magnitude of the coefficient on the interaction term falls somewhat, from -0.095 to -0.071 , when income and population are added (column [2]), but is not further affected when their growth rates are also added (column [3]), and remains significant.

Because housing is a durable good, the argument that we laid out is most appropriate for positive demand shocks. Yet the relationship apparently holds for both positive and negative shocks. Column (4) restricts the sample to observations in which the (real) price index change is positive, which is two-thirds of the sample. The coefficient on the interaction term doubles in magnitude, to -0.15 , and remains significant. Column (5) considers the remaining third of the sample, those with a fall in the price index. The estimated coefficient is, at -0.26 , even larger in magnitude, and significant.¹⁵

3.4 Related Literature

Our work is related to several strands of the housing literature. First, it relates to the literature on MSA-level supply elasticity, which has received increasing attention over the last few years, after having been overshadowed by a much more voluminous literature on demand (e.g., Rosenthal 1999). Recent work has related housing supply elasticity to a number of factors, such as land use regulation (Linneman et al. 1990; Gyourko, Saiz, and Summers 2008) and the share of buildable land (Glaeser, Gyourko, and Saiz 2008; Saiz 2010). In our empirical cross-market analysis, we control for regulation and topography, but highlight the implications of a new supply elasticity proxy—commuting time, which essentially stems from differences in the urban form.

Second, prior research on spatial variation in house price movements has mostly focused on across-market differences. Examples include Glaeser, Gyourko, and Saiz (2008), Van Nieuwerburgh and Weill (2010), and Saiz (2010). As shown in this chapter and elsewhere (e.g., Guerrieri, Hartley, and Hurst 2010), there is significant and systematic within-city variation in house price growth. Thus, both the within-market and cross-market analysis are essential to understanding house price movements. Our work adds to this literature by looking at a single indicator of supply conditions whose variation within- and cross-markets determines the extent of the response of both price and quantity to demand in both contexts.

15. We drop the income and population growth terms in the split sample regressions as they are insignificant. The results are robust to including them or dropping the income and population levels.

Despite its importance, the literature on within-market house price movements is relatively thin. Case and Mayer (1995) and Case and Marynchenko (2002) examine house price movements during the 1980s and early 1990s across zip codes within Boston, Chicago, and Los Angeles. More recently, Landvoigt, Piazzesi, and Schneider (2012) use an assignment model to explain the greater appreciation of low-quality homes in San Diego. Guerrieri, Hartley, and Hurst (2010) explore house price changes across neighborhoods within Chicago, and find that poor areas adjacent to rich areas appreciate more quickly than other areas. Molloy and Shan (2010) use data on a large number of zip codes and municipalities from 1981 to 2008, and find that a 10 percent increase in gas prices leads to a 10 percent decrease in construction in locations with a long average commuting time relative to other locations, but to no significant change in house prices.

While both of these last two papers control for commuting time in the price analysis, the within-market analysis in this chapter differs from their analysis both in its focus and in its implementation. While their papers concentrate on the variation in demand-side determinants (neighborhood gentrification opportunities and the budget share of gas prices), ours is concerned with differing supply elasticities. In the implementation, our work differs in allowing the effect of commuting time to differ across years and MSAs, and, indeed, that is our focus.

As noted earlier, Green, Malpezzi, and Mayo (2005) also derive the supply elasticity from the Capozza and Helsley (1989) model. They estimate a supply elasticity for each MSA by the OLS estimate of an MSA-specific regression of permits per population on average prices, and then regress that estimate on a number of variables, including log population, a per unit distance commuting time (i.e., T), and a land use regulation proxy. Their sample is much smaller than ours, comprising only forty-five MSAs.

3.5 Conclusion

Recent work by Glaeser, Gyourko, and Saiz (2008) has shown, both theoretically and empirically, that house prices increase more during housing booms in places where housing supply is less elastic. In this chapter, we follow their insight and explore the role of a new supply proxy—commuting time—in explaining the within-market and cross-market variation in how house price varies with stages of the market cycle.

In the within-market analysis, we use self-reported home values from the 2005 to 2010 ACS samples to examine the changing relationship between housing price growth and commuting time. Consistent with the notion that building is easier to build at the city edge, we find that the price gradient became flatter in the bust, implying prices fell more in the center than at the city's edge. We do not, however, find any changing relationship over the short part of the boom (2005 to 2006/2007). These results are very similar to the findings on land prices in Haughwout et al.'s chapter 2 in this vol-

ume, although their distance metric—physical distance from the city’s tallest building, grouped and normalized by the distribution of land transactions in each city—is different from ours, and they consider only fifteen large cities. Their table 2.4 shows no difference in the boom years increase in land prices between the middle half and the inner quarter, but a much greater fall for the inner quarter in the bust years.¹⁶

We also find that rents behave qualitatively in the same manner as values. Indeed, we are unable to reject the hypothesis that the price and rent gradients rotate to the exact same degree in the bust. However, the point estimates indicate that the rotation of the rent gradient is more muted than that of prices, which is what we expect from the bursting of an asset bubble, with some overhang effect on rents.

We also consider housing price growth across markets. We argue that the supply elasticity should be smaller in cities with larger average commuting time. Consistent with that, we find that when prices increase overall, prices rise more in MSAs with greater average commuting time, even when controlling for regulatory and topological supply constraint proxies previously used in the literature. We have also checked whether the cross effects of commuting time, and the supply constraints, can explain the finding in Sinai’s chapter 1 in this volume that the cross-sectional dispersion in housing price growth is greater in booms, but we have found no significant relation between the variance and the mean of predicted price growth across years.

As a complementary analysis to that of prices, we also investigate how permits behave over the cycle. Within markets, we find no evidence that the log difference in permits issued in the central and noncentral counties of MSAs differ over that or other cycles. Across markets, however, we find that a given increase in prices is associated with a smaller increase in permits, the greater is average commuting time.

Since Rosenthal’s (1999) lament on the limited work on the supply side of housing, a number of studies have identified regulation and topographical conditions as determinants of supply elasticity. Relying on both economic-spatial reasoning and empirical evidence broadly consistent with it, this chapter has shown commuting time, at both the cross-market and within-market level, to be an important determinant of supply elasticity as well.

References

Anas, A., and H. Rhee. 2007. “When Are Urban Growth Boundaries Not Second-Best Policies to Congestion Tolls?” *Journal of Urban Economics* 61:263–86.

16. The comparison of the middle half to the inner quarter is most akin to our gradient rotation, which compares predicted values at the median and zero commuting minutes. They are able to observe the boom from 2000 on.

- Brueckner, J., and S. Rosenthal. 2009. "Gentrification and Neighborhood Housing Cycles: Will America's Future Downtowns Be Rich?" *The Review of Economics and Statistics* 91 (4): 724–43.
- Bucks, B., and K. Pence. 2008. "Do Borrowers Know Their Mortgage Terms?" *Journal of Urban Economics* 64 (2): 218–33.
- Buonaccorsi. 2010. *Measurement Error: Models, Methods, and Applications*. Boca Raton, FL: Chapman and Hall/CRC.
- Burchfield, M., H. G. Overman, D. Puga, and M. A. Turner. 2006. "Causes of Sprawl: A Portrait from Space." *Quarterly Journal of Economics* 121 (2): 587–633.
- Capozza, D., and R. Helsley. 1989. "The Fundamentals of Land Prices and Urban Growth." *Journal of Urban Economics* 26 (3): 295–306.
- Case, K., and M. Marynchenko. 2002. "Home Appreciation in Low- and Moderate-Income Markets." In *Low Income Homeownership: Examining the Unexamined Goal*, edited by Nicolas Retsinas and Eric Belsky. Cambridge, MA, and Washington, DC: Joint Center for Housing Studies at Harvard University and Brookings Institution Press.
- Case, K., and C. Mayer. 1995. "The Housing Cycle in Eastern Massachusetts: Variation among Cities and Towns." *New England Economic Review* March/April: 24–40.
- Cotter, J., S. Gabriel, and R. Roll. 2011. "Integration and Contagion in US Housing Markets." October 18. Available at Social Science Research Network: <http://ssrn.com/abstract=1945975>.
- DiPasquale, D., and C. T. Somerville. 1995. "Do House Price Indices Based on Transacting Units Represent the Entire Stock? Evidence from the American Housing Survey." *Journal of Housing Economics* 4 (3): 195–229.
- Fuller, W. 1987. *Measurement Error Models*. Hoboken, NJ: John Wiley and Sons.
- Genesove, D. 2003. "Nominal Rigidity in Apartment Rents." *Review of Economics and Statistics* 85 (4): 844–53.
- Glaeser, E., J. Gyourko, and A. Saiz. 2008. "Housing Supply and House Bubbles." *Journal of Urban Economics* 64 (2): 198–217.
- Glaeser, E., and B. Ward. 2009. "The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston." *Journal of Urban Economics* 65 (3): 265–78.
- Green, R., S. Malpezzi, and S. Mayo. 2005. "Metropolitan-Specific Estimates of the Price Elasticity of Supply of Housing and Their Sources." *American Economic Review P&P* 95 (2): 334–39.
- Guerrieri, V., D. Hartley, and E. Hurst. 2010. "Endogenous Gentrification and Housing Price Dynamics." NBER Working Paper no. 16237. Cambridge, MA: National Bureau of Economic Research.
- Gyourko, J., A. Saiz, and A. Summers. 2008. "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies* 45:693.
- Kiel, K., and J. Zabel. 1999. "The Accuracy of Owner-Provided House Values: The 1978–1991 American Housing Survey." *Real Estate Economics* 27 (2): 263–98.
- Landvoigt, T., M. Piazzesi, and M. Schneider. 2012. "The Housing Market(s) of San Diego." NBER Working Paper no. 17723. Cambridge, MA: National Bureau of Economic Research.
- Linneman, P., A. Summers, N. Brooks, and H. Buist. 1990. "The State of Local Growth Management." Harton Real Estate Working Paper no. 81. University of Pennsylvania.
- Mayer, C., and T. Somerville. 2000. "Residential Construction: Using the Urban Growth Model to Estimate Housing Supply." *Journal of Urban Economics* 48 (1): 85–109.

- Molloy, R., and H. Shan. 2010. "The Effects of Gasoline Prices on Household Location." Federal Reserve Board Finance and Economics Discussion Series Working Paper no. 2010-36.
- Rosenthal, S. 1999. "Housing Supply: The Other Half of the Market." *Journal of Real Estate Finance and Economics* 18 (1): 5–8.
- Saiz, A. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3): 1253–96.
- Van Nieuwerburgh, S., and P. Weill. 2010. "Why Has House Price Dispersion Gone up?" *Review of Economic Studies* 77 (4): 1567–606.
- White, M. J. 1988. "Urban Commuting Journeys Are Not 'Wasteful.'" *Journal of Political Economy* 96 (5): 1097–110.