

Endogenous Gentrification and Housing Price Dynamics*

Veronica Guerrieri[†]

University of Chicago and NBER

Daniel Hartley[‡]

Federal Reserve Bank of Cleveland

Erik Hurst[§]

University of Chicago and NBER

August 12, 2011

Abstract

In this paper, we begin by documenting substantial variation in house price growth across neighborhoods within a city during city wide housing price booms. We then present a model which links house price movements across neighborhoods within a city and the gentrification of those neighborhoods in response to a city wide housing demand shock. A key ingredient in our model is a positive neighborhood externality: individuals like to live next to richer neighbors. This generates an equilibrium where households segregate based upon their income. In response to a city wide demand shock, higher income residents will choose to expand their housing by migrating into the poorer neighborhoods that directly abut the initial richer neighborhoods. The in-migration of the richer residents into these border neighborhoods will bid up prices in those neighborhoods causing the original poorer residents to migrate out. We refer to this process as “endogenous gentrification”. Using a variety of data sets and using Bartik variation across cities to identify city level housing demand shocks, we find strong empirical support for the model’s predictions.

*The authors would like to thank seminar participants at Chicago, Cleveland State, Duke Conference on Housing Market Dynamics, Harvard, MIT, Oberlin, Ohio State, Queen’s University Conference on Housing and Real Estate Dynamics, Rochester, Stanford, Summer 2010 NBER PERE meeting, Tufts, UCLA, UIC, University of Akron, Wharton, Winter 2010 NBER EFG program meeting, Wisconsin, and the Federal Reserve Banks of Atlanta, Boston, Chicago, Cleveland, Minneapolis, and St. Louis. We are particularly indebted to Daron Acemoglu, Gary Becker, Hoyt Bleakley, V. V. Chari, Raj Chetty, Morris Davis, Ed Glaeser, Matt Kahn, Larry Katz, Jed Kolko, Guido Lorenzoni, Enrico Moretti, Kevin Murphy, Matt Notowidigdo, John Quigley, Esteban Rossi-Hansberg, Jesse Shapiro, and Todd Sinai for their detailed comments on previous drafts of this paper. All remaining errors are our own. Guerrieri and Hurst would like to acknowledge financial support from the University of Chicago’s Booth School of Business. The views expressed herein are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

[†]Contact: *Veronica.Guerrieri@chicagobooth.edu*

[‡]Contact: *Daniel.Hartley@clev.frb.org*

[§]Contact: *Erik.Hurst@chicagobooth.edu*

1 Introduction

It has been well documented that there are large differences in house price appreciation rates across U.S. metropolitan areas.¹ For example, according to the Case-Shiller Price Index, real property prices increased by over 100 percent in Washington DC, Miami, and Los Angeles between 2000 and 2006, while property prices appreciated by roughly 10 percent in Atlanta and Denver during the same time period. Across the 20 MSAs for which a Case-Shiller MSA index is publicly available, the standard deviation in real house price growth between 2000 and 2006 was 42 percent. Such variation is not a recent phenomenon. During the 1990s, the Case-Shiller cross-MSA standard deviation in house price growth was 21 percent.

While most of the literature has focused on trying to explain cross-city differences in house price appreciation, we document that there are also substantial *within*-city differences in house price appreciation. For example, between 2000 and 2006 residential properties in the Harlem neighborhood of New York City appreciated by over 130 percent, while residential properties less than two miles away, in midtown Manhattan, only appreciated by 45 percent. The New York City MSA, as a whole, appreciated by roughly 80 percent during this time period. Such patterns are common in many cities. Using within-city price indices from a variety of sources, we show that the average *within*-MSA standard deviation in house price growth during the 2000 - 2006 period was roughly 20 percent. Similar patterns are also found during the 1990s and 1980s. As is commonly discussed in the popular press, these large relative movements in property prices within a city during city-wide property price booms are often associated with changing neighborhood composition. Returning to the Harlem example, a recent New York Times article discussed how Harlem residents have gotten richer during the period when its house prices were substantially appreciating.²

Our goals in this paper are threefold. First, we set out to document a new set of facts about the extent and nature of within-city house price movements during citywide housing price booms. The house price appreciation for the city as a whole is just a composite of the house price movements within all the neighborhoods of the city. Therefore, understanding the movements in house prices across neighborhoods within a city is essential for understanding house price movements for the entire city. Using a variety of different data sources, we show that there are substantial differences across neighborhoods within a city with respect to their house price growth when the city as a whole experiences a housing price boom. Moreover, we show that there is a systematic pattern in this variation. In particular, we document three facts that are robust across time and data sources with respect to within-city house price movements. First, during city-wide housing price booms, neighborhoods with low initial housing prices appreciate at much greater rates than neighborhoods with high initial prices. Second, the variation in housing price appreciation rates among low housing price neighborhoods is much higher than the variation in housing price appreciation rates for higher housing price neighborhoods. Finally, we show that the larger the

¹See, for example, Davis et al. (2007), Glaeser et al. (2008), Van Nieuwerburgh and Weill (2010), and Saiz (2010).

²See the article "No Longer Majority Black, Harlem Is In Transition" from the January 5th, 2010 New York Times.

city-wide housing price boom, the greater is the difference in housing price appreciation rates between low house price and high house price neighborhoods. Regardless of the interpretation we give to some of these facts in later sections, we feel these facts alone are an interesting contribution to the literature on spatial variation in housing price growth.

Our second goal is to develop a spatial model of a city that links within-city neighborhood housing price dynamics with within-city neighborhood gentrification. We represent a city as the real line and each point on the line is a location. Agents are fully mobile across location and there is a representative firm that can build houses in any location at a fixed marginal cost. The key ingredients of the model are that agents are heterogeneous in their income and all agents prefer to live close to richer neighbors. The relevance of such a neighborhood consumption externality in determining house prices is supported by the recent empirical work of Bayer et al. (2007) and Rossi-Hansberg et al. (2010). We show that there exists an equilibrium with full income segregation where the high income residents are concentrated all together and the low income residents live at the periphery. The sorting, as in Becker and Murphy (2003), is the result of the neighborhood externality where all agents are willing to pay more to live closer to rich neighbors. Given diminishing marginal utility, poor residents, however, are less willing to pay high rents to live in the rich neighborhoods, so in equilibrium they live farther from the rich. Within the model, house prices achieve their maximum in the richer neighborhoods and decline as one moves away from them, to compensate for the lower level of the externality. For the neighborhoods that are far enough from the rich, there is no externality, and house prices are equal to the marginal cost of construction.

One of the main contribution of our model, and the basis for our subsequent empirical work, is to explore the dynamics of house prices across neighborhoods in response to city-wide housing demand shocks. Although there is no supply constraint and the city can freely expand, average house prices increase in response to an increase in city-wide housing demand because of gentrification. In particular, the neighborhoods that endogenously getrify are the poor neighborhoods on the border of rich neighborhoods. For concreteness, we say that a neighborhood gentrifies when some poor residents are replaced by richer ones, increasing the extent of the neighborhood externality.³ For example, we consider a city hit by an increase in labor demand and a subsequent wave of migration (Blanchard and Katz, 1992). The richer migrants prefer to locate next to the existing richer households. As a result, they bid up the land prices in the poor neighborhoods that are next to the rich neighborhoods causing the existing poor residents to move out and the city as a whole to expand. Gentrification is the endogenous response to the city-wide housing demand shock and the gentrifying neighborhoods are the poor neighborhoods on the boundry of the richer neighborhoods that experience the largest housing price increase.

³McKinnish et al. (2010) define gentrifying neighborhoods as poor neighborhoods that experience an increase in average income above a certain threshold over a specific period of time, which is slightly different but consistent with our definition. In Section 7, we discuss some of the existing literature on gentrification in much greater depth. Similarly, Kolko (2007) defines gentrification as any neighborhood with positive income growth in an initially lower-income, central-city census tract. Vigdor et al. (2002) uses a definition of gentrification that involves entry of new residents that have higher socioeconomic status than the current residents which may or may not displace the original residents.

Our model predicts that, in response to a positive city-wide housing demand shock, land prices in some poor neighborhoods appreciate at a much faster rate than both richer neighborhoods and other poorer neighborhoods. In particular, the poor neighborhoods that are in close proximity to the rich neighborhoods are the ones that have housing prices that increase the most. Also, our mechanism implies that unexpected permanent shocks to housing demand lead to permanent increases in house prices at the city level although the size of the city is completely elastic. This happens because gentrification bids up the value of the land in the gentrifying neighborhoods. Given this, we find that average price growth within the city is affected both by the size of the housing demand shock and by the particular shape of preferences, technology, and income distribution within the city.

Our third goal is to provide explicit evidence showing that our endogenous gentrification mechanism is an important determinant of within-city variation in house price growth in response to city-wide housing demand shocks. We do this in multiple ways. To begin, we provide an additional fact about within-city neighborhood house price appreciation during city-wide housing booms. In particular, we show that, as our theory predicts, the poor neighborhoods that are next to the rich neighborhoods are the ones that appreciate the most during city-wide housing booms. This result holds in the 1980s, 1990s, and 2000s and holds using a variety of different measures of neighborhood housing price appreciation. Moreover, these results are robust to including controls for distance to the city's center business district, the average commuting time of neighborhood residents, and proximity of the neighborhood to fixed natural amenities such lakes, oceans, and rivers. Again, these results are consistent with the first order predictions of our model.

We then proceed to provide empirical tests that are more directly linked to the core mechanism in our model. Using a Bartik-style instrument to isolate exogenous city level housing demand shocks (Bartik, 1991), we show that it is the housing prices in poor neighborhoods next to rich neighborhoods that appreciate the most in response to the exogenous city-wide housing demand shocks. Our Bartik shock predicts expected income growth in a city between periods t and $t+k$ based on the initial industry mix in that city at time t and the change in industry earnings for the entire U.S. between t and $t+k$. For example, in response to a one standard deviation Bartik shock, poor neighborhoods within the city which directly border a rich neighborhood have housing prices that appreciate roughly 6.8% more than otherwise similar poor neighborhoods within the city that are more than 3 miles away from rich neighborhoods. Again, these results hold controlling for distance to the center business district and proximity to fixed natural amenities within the city.

Finally, we explicitly show that the neighborhoods that appreciate the most during the exogenous city-wide housing demand shock also gentrify. Gentrification - the out migration of poor residents and the in migration of rich residents - is the key mechanism for the within-city house price dynamics we highlight. For this analysis, we again explore the within-city response to a Bartik-style shock. In particular, we show

that in response to an exogenous city-wide demand shock, poor neighborhoods close to rich neighborhoods experience larger increases in neighborhood income, larger increases in the educational attainment of neighborhood residents, and larger declines in the neighborhood poverty rate than do otherwise similar poor neighborhoods that are farther away from the rich neighborhoods. For example, average neighborhood income grows by roughly 11 percent more in response to a one standard deviation Bartik shock for poor neighborhoods that border the rich neighborhoods than it does for otherwise similar poor neighborhoods that are more than 3 miles away from the rich neighborhoods. Lastly, we highlight that during both the 1980s and 1990s, almost all the poor neighborhoods that did in fact gentrify by some ex-post criteria were neighborhoods that were directly bordering existing rich neighborhoods.

In the last section of our paper, we place our paper in context by discussing how our results add to the existing literatures on within-city house price dynamics, urban gentrification, neighborhood consumption externalities, Tiebout sorting, and residential segregation. In addition, we discuss some of the outstanding issues with respect to our empirical methodology including the focus on housing prices instead of either land prices or rental rates, the potential for mean reversion in land/housing prices, and the role of expectations and uncertainty. Finally, we discuss how our results on within-city housing price dynamics can inform our understanding of cross-city housing price dynamics.

Before proceeding, we would like to make two additional comments with respect to our work. First, as noted above, a key ingredient in our model is the existence of neighborhood consumption externalities in that individuals get utility from having rich neighbors relative to poor neighbors. Although, we do not explicitly model the direct mechanism for the externality, we have many potential channels in mind. For example, crime rates are lower in richer neighborhoods. If households value low crime, individuals will prefer to live in wealthier neighborhoods. Likewise, the quality and extent of public goods may be correlated with the income of neighborhood residents. For example, school quality - via peer effects, parental monitoring, or direct expenditures - tends to increase with neighborhood income. Finally, if there are increasing returns to scale in the production of desired neighborhood amenities (number and variety of restaurants, easier access to service industries such as dry cleaners, movie theaters, etc.), such amenities will be more common as the income of one's neighbors increases. Although we do not take a stand on which mechanism is driving the externality, our preference structure is general enough to allow for any story that results in higher amenities being endogenously provided in higher income neighborhoods.

Second, although we prefer to highlight the existence of neighborhood consumption externalities to link within-city house price appreciation and neighborhood gentrification, other traditional urban stories could yield similar theoretical results. For example, if cities are viewed as centers of production agglomeration, as in the classic work by Alonso (1964), Mills (1967), and Muth (1969), neighborhoods that are close to jobs will have higher land prices than neighborhoods that are farther away. Additionally, the models put forth by Rosen (1979) and Roback (1982) show that land prices within the city can differ based on

their proximity to a desirable fixed natural amenity (like proximity to the ocean or beautiful vistas).⁴ Depending on the nature of household preferences, both of these types of models could also generate the link we emphasize between neighborhood gentrification and housing price dynamics in response to a city-wide housing demand shock. In this paper, we focus on the consumption externality story as opposed to the other traditional urban stories. As different types of residents move in and out of the neighborhoods, the amenities those neighborhoods provide endogenously change. Even though the other urban stories can theoretically generate similar patterns, empirically they do not seem to drive the relationships we observe in the data. We control directly for proximity to jobs and proximity to fixed natural amenities in our empirical work. As we show consistently, it is proximity to rich neighborhoods that seems to determine the gentrification patterns we document above and beyond the proximity to jobs and proximity to fixed natural amenities.

2 New Facts About Within City Housing Price Dynamics

In this section, we outline a series of new facts about the nature of housing price dynamics across different neighborhoods within a city (MSA) during city (MSA) wide housing price booms.⁵ Unlike previous attempts to study within-city house price movements, we analyze the patterns simultaneously for a large number of cities and for multiple time periods.⁶ As we show, there are many systematic patterns that emerge with respect to house price dynamics across neighborhoods within a city during city-wide housing price booms.

2.1 House Price Data

Throughout the paper, we primarily use three separate data series to examine within-city house price movements. Each of the series has different strengths and weaknesses. However, despite the differences, the empirical results we emphasize can be found using all three housing price series.

Our primary measure of within-city house price growth comes from the Case-Shiller zip code level price indices.⁷ The Case-Shiller indices are calculated from data on repeat sales of pre-existing single-family homes. The benefit of the Case-Shiller index is that it provides consistent constant-quality price indices for localized areas within a city or metropolitan area over long periods of time. Most of the Case-Shiller zip

⁴The Rosen (1979) and Roback (1982) models were built to explain cross-city variation in housing prices but can be naturally extended to explain within-city variation in housing prices.

⁵Throughout this section, we will often use the term city and MSA interchangeably in our discussion. However, for each of our empirical results, we will be explicit about whether we are exploring within-city or within-MSA dynamics. In our empirical work in Sections 4 and 5, which tests for the importance of endogenous gentrification in explaining within-area house price dynamics, we will focus our results on variation within cities. Doing so allows us to hold factors that could conceivably vary across cities (like tax rates and public expenditures) fixed.

⁶See, for example, Case and Mayer (1996) and Case and Marynchenko (2002).

⁷The zip code indices are not publicly available. Fiserv, the company overseeing the Case-Shiller index, provided them to us for the purpose of this research project. The data are the same as the data provided to other researchers studying local movements in housing prices. See, for example, Mian and Sufi (2009). Unfortunately, we only have the data through 2008 and, as a result, we cannot systematically explore within-city house price patterns during the recent bust. We have been unsuccessful in our attempts to secure the post 2008 data from Fiserv.

code-level price indices go back in time through the late 1980s or the early 1990s. The data was provided to us at the quarterly frequency and the most recent data we have access to is for the fourth quarter of 2008. As a result, for each metro area, we have quarterly price indices on selected zip codes within selected metropolitan areas going back roughly 20 years.

There are a few things that we would like to point out about the Case-Shiller indices. First, the Case-Shiller zip code level indices are only available for certain zip codes in certain metropolitan areas. For some of our analysis, we focus our attention only on the zip codes within the main city in the MSA. For example, we look at the patterns within the city of Chicago instead of just the broad Chicago MSA. When doing so, we only use the MSAs where the main city within the MSA has at least 10 zip codes with a usable house price index.⁸

Second, we only use information for the zip codes where the price indices were computed using actual transaction data for properties within the zip code. Some of the zip code price indices computed as part of the available Case-Shiller data use imputed data or data from some of the surrounding zip codes. We exclude all such zip codes from our analysis. As a result, the Case-Shiller zip codes that we use in our analysis do not cover the universe of zip codes within a city. Only about 50 percent of the zip codes in the city of Chicago, for example, have housing price indices computed using actual transaction data. The fraction in other cities is closer to 100 percent. The zip codes within the cities that tend to have either missing or imputed zip code housing price indices are the zip codes where there are very few housing transactions or where most of the housing transactions are not for single-family homes. Restricting our analysis to the primary Case-Shiller cities (within each metro area) and to the zip codes with price indices based on actual transaction data, we have data for 508 zip codes during the 1990-2000 period and for 497 zip codes during the 2000-2006 period. If we expand our analysis to the entire metro area where we have a housing price index based on non-imputed data, we have 1,529 zip codes during the 1990-2000 period and 1,693 zip codes during the 2000-2006 period.

Third, the Case-Shiller index has the goal of measuring the change in land prices by removing structure fixed effects using their repeat sales methodology. However, this methodology only uncovers changes in land prices if the attributes of the structure remain fixed over time. If households change the attributes of the structure via remodeling or through renovations, the change in the house prices uncovered by a repeat sales index will be a composite of changes in land prices and of improvements to the housing structure. Those who compute the Case-Shiller index are aware of such problems and, albeit imperfectly, take steps to minimize the effect of potential remodeling and renovations.⁹

⁸According to this criterion, we focus on 26 cities: Akron, Atlanta, Charlotte, Chicago, Cincinnati, Columbus (OH), Denver, Fresno, Jacksonville, Las Vegas, Los Angeles, Memphis, Miami, New York, Oakland, Philadelphia, Phoenix, Portland (OR), Raleigh, Sacramento, San Diego, San Jose, Seattle, St. Paul, Tampa, and Toledo. Any analysis at the zip code level in which we broaden our sample to examine the full MSA uses the same 26 MSAs. Boston, San Francisco, and Washington D.C., which have less than 10 zip codes with usable price indices within the main city, are not included in the within city analysis but are also included when we perform our within-MSA analysis.

⁹In particular, the index puts a lower weight on repeated sales transactions where the change in price is likely to reflect

Given the discussion above, we see four main limitations to the Case-Shiller index. First, the Case-Shiller index is only available for the MSAs selected to be part of the Case-Shiller series (which tend to be the larger MSAs). Second, the series is not always available for all zip codes within the Case-Shiller MSAs. Third, the index may not be perfectly capturing changes in land prices because it cannot perfectly control for unobserved renovations or remodeling. Finally, the lowest level of aggregation available for the Case-Shiller data is at the level of the zip code.

We take all four of the above concerns seriously and try to address them by augmenting our analysis with data from two other sources. First, we use zip code level data on house prices from the Zillow Home Value Index.¹⁰ Instead of using a repeat sales methodology, Zillow uses the same underlying deed data as the Case-Shiller index but creates a hedonically adjusted price index. The Zillow index uses detailed information about the property, collected from public records, including the size of the house, the number of bedrooms, and the number of bathrooms. To the extent that the average measured characteristics of the home change over time, the Zillow index will capture such changes. The Zillow data is available at the monthly level for most zip codes within the U.S. starting in the late 1990s. Finally, even in the Zillow data, some zip codes do not have enough transactions during the month to create a reliable house price index. The Zillow data that we have access to indicates the zip codes that Zillow feels do not have enough transactions to create a reliable price index. We exclude such zip codes from our analysis.

The Zillow index, at least partially, overcomes some of the deficiencies of the Case-Shiller index in that it potentially allows for the broad attributes of the structure (e.g., square footage, number of bedrooms, etc.) to change over time. Also, a reliable Zillow index is available for more zip codes within a city than the Case-Shiller index. The reason for this is that the Case-Shiller index is based off of repeat sales transactions while the Zillow index uses all sales regardless of whether or not they could match the sale with a previous transaction. Finally, the Zillow results are available for more cities. The Zillow data allows us to see if our results using the Case-Shiller data change in any substantial way when we include a broader set of zip codes.¹¹ Neither the restricted set of zip codes nor the failure to control for changing structure attributes modify any of our key empirical results in any way.

Finally, we augment our results using information on the percent change in median house price at

changes in the housing structure, that is, when the change in price was either disproportionately large or disproportionately small. Additionally, the index excludes all properties where the property type changed (i.e., a single family home is converted to condos) and it excludes all properties where the home sells within six months after a purchase. These properties tend to follow the redevelopment of the property. Also, all repeated sales transactions are weighted based on the time interval between first and second sales. Sales pairs with longer time intervals are given less weight than sales pairs with shorter intervals. The assumption is that if a sales pair interval is longer, then it is more likely that a house may have experienced a physical change. For more information on the construction of the Case-Shiller indices see the Standard and Poor's web-site which documents their home price index construction methodology. See <http://www.caseshiller.fiserv.com/about-fiserv-case-shiller-indexes.aspx>.

¹⁰See <http://www.zillow.com/wikipages/Real-Estate-Market-Reports-FAQ/> for details. We thank Amir Sufi for providing us with the zip code level indices including the information he received from Zillow on which zip codes had too few observations to make a reliable price index. We posted all such data on our web pages. See the online robustness appendix for details.

¹¹There are a few instances where there is an available Case-Shiller house price index for a zip code but there is not a corresponding Zillow index. As a result, our actual sample size in some specifications where we restrict the sample to zip codes where both indices exist is slightly lower than the sample sizes for the Case-Shiller samples discussed above.

the neighborhood level from the 1980, 1990, and 2000 U.S. Censuses. The primary benefit of the Census data is that it is available at very fine levels of spatial aggregation.¹² In particular, we can examine within-city differences in housing price dynamics at both the level of zip codes and census tracts. We compute within-zip code or within-census tract appreciation rates by computing the growth in the median house price across similarly defined levels of disaggregation between 1980 and 1990 and between 1990 and 2000. The Census data, however, are not without limitations. Unlike the repeat sales methodology of the Case-Shiller index or the hedonic method of the Zillow index, the Census data is simply the growth in the median house price within a zip code or census tract. As a result, it may be confounding movements in land prices with movements in structure quality for the median house. Moreover, the median house value, in terms of quality, could be changing over time. For example, as low quality housing gets demolished, the median price in a neighborhood may increase with no change in either land prices or structure attributes for the remaining properties. We can partially address this limitation by including controls for the changes in neighborhood housing stock characteristics when using this measure. As we show below, the patterns of zip code level house price appreciation found with the Census data are quite similar to the patterns found with the other two data sources. As a result, we feel confident in using the Census data to explore house price dynamics at the sub-zip code level.

In the Appendix (Tables A1 and A2), we show that the different house price series do, in fact, track each other quite closely for the zip codes where multiple indices exist.¹³ For example, between 2000 and 2006, the zip code level house price appreciation rates from the Case-Shiller data and from the Zillow data are correlated at between 0.95 and 0.96 depending on whether we look at all available zip codes within the MSA or if we restrict the sample to only the zip codes within the main city of the MSA. For the 1990-2000 period, the correlation between the Case-Shiller zip code index and the index formed using the Census zip code level data was about 0.8.

In Table A2 we report the regression of the Case-Shiller zip code level appreciation rate on the Zillow zip code level appreciation rate (columns 1 and 2) and then separately on the Census zip code median house price appreciation rate (columns 3 and 4). The first and third columns use the sample with all Case-Shiller zip codes in the MSA. The second and fourth columns restrict the analysis to the zip codes in the main city of the MSA. The purpose of this table is to see if the relationship between the appreciation rates of the various data sources diverges at higher levels of appreciation. The answer appears to be no. The coefficient on the house price growth within the main cities is essentially one for both the Zillow data (column 2) and for the Census data (column 4). It should be noted that the levels of the price appreciations do differ across the surveys (as seen by the constant estimates in Appendix Table A2). But, given that most of our results are going to be identified off of variation in house price appreciation across

¹²We thank Ed Glaeser for the suggestion of moving some of our analysis to the sub zip code level using the Census data.

¹³All data in the paper are reported in real 2000 prices, unless otherwise indicated. Likewise, all growth rates are in real terms. We use the CPI-U (all items less shelter) to convert the nominal variables into real variables.

neighborhoods within a city, differences in the level of the price appreciation, either across cities or across price indices for a given city, do not alter our results. The results in Appendix A1 and A2 illustrate that the house price appreciation rates across the various data sources are highly correlated. As a result, it is not surprising that the results we document below are very similar regardless of the house price measures we use.

Throughout the paper, we compute MSA level house price appreciation rates using Federal Housing Finance Agency (FHFA) metro level housing price indices if the Case-Shiller house price series is not available for the MSA. For the MSAs where both data sets exist, the Case-Shiller and FHFA data track each other nearly identically.¹⁴

Finally, for some of our descriptive results in this section (and for the Harlem example in the introduction), we use detailed data on New York City neighborhood home price appreciation from the Furman Center repeat sales index which covers all neighborhoods within New York City. The Furman Center data use NYC community districts as their level of aggregation as opposed to zip codes and, as a result, have enough observations to make reliable indices for all areas within NYC. The Case-Shiller and Zillow indices do not cover much of the zip codes of New York City proper (although they do provide indices for many zip codes in the New York MSA).¹⁵

2.2 Within City Housing Price Dynamics

While much work has documented the variation in house prices across cities, little work has been done to systematically document the variation in house prices within cities. In this subsection, we document three new facts about within-city house price movements. After presenting the model of endogenous gentrification in Section 3, we will revisit these facts to both interpret them and put more structure on them. However, we view these facts as being important not only for the model we are trying to highlight, but are also of independent interest.

Table 1 shows the degree of between- and within-MSA variation in house price appreciation separately during the 2000-2006 period (row 1), the 1990-2000 period (row 2) and the 1980-1990 period (row 3). Columns 1 and 2 focus on cross-MSA variation in house price appreciation for comparison to the within-MSA or within-city variation. When focusing on the cross-MSA variation, we use data from two sources. Column 1 uses data from the FHFA MSA level house price indices, while Column 2 uses data from the Case-Shiller MSA level house price indices. The Case-Shiller MSA level house price appreciation is only available for a handful of MSAs during the 1980s, so we only provide the FHFA results for this time period. For reference, the house price appreciation rates using the Case-Shiller MSA level index for the 1990-2000

¹⁴It has been shown that the national Case-Shiller appreciation rates and the FHFA appreciation do not track each other, particularly in recent years. However, this is entirely due to differences in the regions covered by the two surveys. For MSAs where both data series exist, the trends in appreciation rates are nearly identical even in recent periods.

¹⁵See <http://furmancenter.org/>. There are 59 community districts in NYC which represent clusters of several zip codes. The Furman data for NYC extend back to 1974.

period and for the 2000-2006 period for each MSA are shown in Appendix Table A3. As seen from Table 1 and Appendix Table A3, there is large variation in price appreciation across MSAs during the 1980s, the 1990s and the 2000s. This is consistent with the well documented facts discussed in Davis et al. (2007), Glaeser et al. (2008), Van Nieuwerburgh and Weill (2010), and Saiz (2010).

The next four columns of Table 1 show within-city or within-MSA, cross-zip code variation in house price appreciation for the same time periods. For columns 3 and 4, we use data from the Case-Shiller indices and show the results for all available zip codes within the MSA (column 3) and then for all available zip codes in the main city of the MSA (column 4). In column 5, we show the within-city cross-zip code standard deviation in house price appreciation using the Zillow data, which is only available for 2000-2006. In column 6, we show the results for the within-city cross-zip code standard deviation in house price appreciation using the Census data for the 1990-2000 period. For columns 5 and 6, we restrict the Zillow sample and the Census sample to match exactly the Case-Shiller sample. This was done so that the results can be compared across the different house price measures. The final two columns show within-city cross-census tract variation for 1990-2000 and 1980-1990 periods. The sample in column 7 is restricted to census tracts that overlap with the 496 zip codes used in the sample for column 6. As might be expected if there is within-zip code variation, growth rates of census tracts show more within-city variation than growth rates of zip codes. Finally, column 8, broadens the sample of census tracts to include all tracts in all cities that contain at least 30 census tracts in the initial period of the sample. This larger sample shows even more within-city variation in housing price growth rates.

One of our key descriptive results are shown in Table 1. Table 1 shows that there is substantial variation in house price appreciation rates across zip codes and census tracts within a city. For example, during the 2000-2006 period, the within-city variation was about one half as large as the cross-city variation but was still substantial, between 18 and 24 percentage points depending on the housing price series. During the 1990-2000 period, the within-city variation was of the same order of magnitude as the cross-city variation. As one would expect, the within-city variation increases as the level of our definition of a neighborhood gets smaller. For example, the cross-census tract variation in house price growth both in the 1980s and the 1990s was roughly fifty percentage points. These results show that within-city variation in house price growth is of at least the same order of magnitude as the cross-city variation that has received so much attention in the literature.

We now highlight that there is some systematic variation in the differential house price growth across neighborhoods within a city. The next fact we wish to highlight is shown in Figure 1 and Table 2. Figure 1 plots the house price appreciation rate in each zip code within the New York MSA between 2000 and 2006 (using the Case-Shiller data) against the median house price for the same zip codes in year 2000 (from the Census). As seen from the figure, there is a sharp negative relationship between the initial level of housing prices within the zip code and the subsequent appreciation rate in the zip code. On average,

zip codes with lower initial housing prices appreciated at roughly twice the rate as zip codes with higher initial housing prices during this period.

Our choice of showing New York in Figure 1 is done for illustrative purposes. Table 2 shows the relationship between the initial median housing price and the subsequent housing price growth across neighborhoods within the city/MSA for a large selection of cities and metro areas during different time periods. Specifically, Table 2 shows the mean growth rate in property prices over the indicated time period for neighborhoods in different quartiles of the initial house price distribution within the city or metro area. The last column shows the p -value of the difference in house price appreciation rates between the properties that were initially in the top (column 1) and bottom (column 4) quartiles of the housing price distribution within the city or metro area. This table is the analog to the scatter plot shown in Figure 1. In all cases, the initial level of housing prices used to define the quartiles in period t is defined using the median level of reported house price for the neighborhood from the corresponding U.S. Census (i.e., 2000, 1990, or 1980 depending on the time period studied). The house price appreciation is measured using the Case-Shiller index unless noted otherwise on the table.

We can conclude a few things from the results in Table 2. First, the patterns found in Figure 1 for New York for the 2000-2006 period are also found in a wide variety of other cities and MSAs during the same period. For example, consider Chicago. Within both the Chicago city and the Chicago MSA, low initial price neighborhoods (quartile 4) appreciated at much higher rates than high initial price neighborhoods (quartile 1) during the 2000-2006 period. Specifically, within the city of Chicago, low price neighborhoods appreciated at close to 90 percent while high price neighborhoods appreciated at about half that rate (50 percent). Also, as seen from Table 2, these within-city patterns are not limited to the recent period. During the 1990s, Denver and Portland experienced large housing price booms, and it was the low priced neighborhoods that appreciated at much higher rates than the high priced neighborhoods. Additionally, during the 1980s, New York and Boston experienced large housing price booms. Also in these cases, it was also the low priced neighborhoods that appreciated at much higher rates than the high priced neighborhoods during this time period. These results suggest that our findings are not specific to the recent housing price boom.

Additionally, there is also some evidence that poor neighborhoods fall the most during city-wide housing price busts. For example, within San Francisco and Boston during the 1990s, the poorer neighborhoods contracted slightly more relative to the richer neighborhoods. In Boston during the 1990s, it was actually the neighborhoods within the third quartile of the initial house price distribution that contracted the most. Given the data we have available, we cannot systematically explore the behavior of house price movements across neighborhoods within a city during city-wide housing price busts. As more data gets released from the recent time period, such an analysis will be possible. However, from the little data we have from the 1990s, it looks like high priced neighborhoods are the least price elastic during both housing price booms

and housing price busts.

Are the results shown in Table 2 representative of the patterns in a broader sample of cities? The answer is definitely yes. To illustrate this, we estimate:

$$\frac{\Delta P_{t,t+k}^{i,j}}{P_t^{i,j}} = \mu_j + \omega_1 \ln(HP_t^{i,j}) + \epsilon_{t,t+k}^{i,j} \quad (1)$$

where $\Delta P_{t,t+k}^{i,j}/P_t^{i,j}$ is the growth in housing prices between period t and $t+k$ within neighborhood i in city or MSA j using the various house price series and $HP_t^{i,j}$ is the median house price in neighborhood i in city or MSA j in year t as measured by the U.S. Census. Given that we also include city or MSA fixed effects, μ_j , all of our identification comes from variation across neighborhoods within a city/MSA. The variable of interest from this regression is ω_1 which estimates the relationship between initial median house prices in the neighborhood and subsequent neighborhood housing price growth. We run this regression using different neighborhood house price series and for different time periods. For all specifications, we weight the data using the number of owner occupied housing units in the neighborhood during period t (from the Census).

To conserve space, we do not show the results of this regression in the main text. However, in the online robustness appendix that accompanies this paper, we show the results of this specification for different time periods, different measures of house price growth, and for different samples. The results across the different specifications are very consistent. For cities experiencing a city wide housing price boom, it is the neighborhoods with the initially low housing prices that appreciate the most. For example, during the 2000-2006 period, restricting the sample to all zip codes with a Case-Shiller house price index, and using the Case-Shiller index to measure zip code housing price growth, our estimate of ω_1 is -0.23 with a standard error of 0.05.

In the robustness appendix, we also formally document another fact about within city house price dynamics. In particular, we show that the difference between the house price appreciation of initially low price neighborhoods within the city and initial high price neighborhoods within the city increases with the size of the city wide housing price boom. For example, during the 2000-2006 period there was no difference in the house price appreciation rates of initially high and initially low price neighborhoods within Columbus, Ohio. However, as shown above, the appreciation rate of initially low price neighborhoods in New York was twice as high as the appreciation rate for initially high price neighborhoods in New York. These patterns also held for the 1990s as well. New York is not an outlier with respect to these patterns. Systematically, the gap between the price appreciation of low price neighborhoods within a city and high price neighborhoods within a city grows as the size of the city wide house price boom increases. We show these patterns hold in the 1990s as well. For cities that experienced slightly declining house values during the 1990s, there was little difference in house price appreciation rates across initially rich and poor neighborhoods. However, for cities like Portland and Denver that experienced large housing price increases

during the 1990s, initially low price neighborhoods appreciated at much higher rates than initially high price neighborhoods. Again, we show the regressions summarizing these results in the online robustness appendix that accompanies the paper.

Finally, and most important for our model and empirical work below, we highlight a third fact about within city house price appreciations during city wide housing price booms. Returning to Figure 1, another feature of the data for the New York MSA is that the house price appreciation rate among initially low priced neighborhoods exhibits substantially more variability than the house price appreciation rates among initially high priced neighborhoods. In particular, the standard deviation of housing price growth between 2000 and 2006 for neighborhoods in the lowest initial house price quartile for the New York MSA was 29 percent while the standard deviation of house price growth between the same time period for neighborhoods in the top initial house price quartile for the New York MSA was only 5 percent. The difference is significant at the less than 1 percent level.

This difference in variability of growth rates between initially low priced neighborhoods and initially high priced neighborhoods within a city during a city-wide housing price boom is a robust feature of the data across the many cities in our sample. Again, we formally document these facts in the online robustness appendix that accompanies the paper. When cities experience housing price booms, the variability in house price growth among initially low price neighborhoods is much higher than the variability of house price growth among initially high price neighborhoods. In particular, some low price neighborhoods appreciate at a rate similar to high price neighborhoods while other low price neighborhoods appreciate at rates much higher than the high price neighborhoods.

In summary, we document three novel facts about the extent of within city housing price growth during city wide house price booms. First, we show that during city-wide housing price booms, neighborhoods with lower initial housing prices appreciated at much higher rates than neighborhoods with higher initial housing prices. Second, we show that the difference between low and high price neighborhood house price growth grows with the size of the city wide housing price boom. Finally, we show that the variation in house price growth among initially low price neighborhoods is much higher than the variation in housing price growth among high price neighborhoods. It is this variation among low priced neighborhoods that we will exploit to directly test the mechanism at the heart of the model we present in the next section. Why is it that some low price neighborhoods within a city (like the Harlem neighborhood in New York City during the 2000s) appreciate at very high rates while other low price neighborhoods (like the Brownsville or Jamaica neighborhoods in New York City during the 2000s) experience much more modest house price growth? Our endogenous gentrification mechanism can explain such variation.

3 Model

In this section, we develop a spatial model of housing prices across neighborhoods within a city so as to explore the relationship between gentrification and house price dynamics in response to city-wide housing demand shocks. The key ingredient of our model is a positive neighborhood externality: people like to live next to richer neighbors. We do not micro-found the source of this externality and leave the model flexible enough to encompass alternative possible stories behind the preference for richer neighborhoods, such as lower crime rates, higher school quality, and more positive neighborhood amenities. Whatever micro-foundation one prefers, the presence of such an externality generates a gentrification process in response to a city-wide increase in housing demand.

Let us mention that similar theoretical results could be obtained in models of a city without neighborhood externalities, where neighborhoods are heterogeneous because of commuting costs and proximity to fixed natural amenities. We choose to focus on a neighborhood consumption externality for two reasons. First, recent empirical works by Bayer et al. (2007) and Rossi-Hansberg et al. (2010) have shown that such a mechanism is important to explain within-city housing price dynamics. Second, in our empirical work below, we explicitly control for proximity to jobs and proximity to fixed natural amenities and show that such controls have little effects on our empirical results. Although we think that these alternative stories may be important in explaining within-city house price differences, we do not think that these mechanisms are at the heart of the link between neighborhood gentrification and neighborhood house price dynamics that we document below. For this reason, in the model we abstract from these other mechanisms.

3.1 Set up

Time is discrete and runs forever. We consider a city populated by N infinitely lived individuals comprised of two types: a continuum of rich households of measure N^R and a continuum of poor households of measure N^P . Each period households of type s , for $s = R, P$, receive an exogenous endowment of consumption goods equal to y^s , with $y^P < y^R$.

The city is represented by the real line and each point on the line $i \in (-\infty, +\infty)$ is a different location.¹⁶ Agents are fully mobile and can choose to live in any location i . Denote by $n_t^s(i)$ the measure of households of type s who live in location i at time t and by $h_t^s(i)$ the size of the house they choose. In each location, there is a maximum space that can be occupied by houses which is normalized to 1,¹⁷ that

¹⁶We choose to model the city as a line because it simplifies our analysis. The main implications of our model extend to a circular city as in Lucas and Rossi-Hansberg (2002).

¹⁷Our notion of space is uni-dimensional: if there is need for more space to construct houses we assume that the neighborhoods have to expand horizontally. We could enrich the model with a bi-dimensional notion of space, by allowing a more flexible space constraint in each location. For example, we could imagine some form of adjustment cost to construct in each location, so that in reaction to a demand shock the city can expand both in the horizontal and in the vertical dimension. Our model is the extreme case with infinite adjustment cost on the vertical dimension and no adjustment costs on the horizontal dimension. Our mechanism would go through if we allow some convex adjustment costs to the vertical margin.

is,

$$n_t^R(i) h_t^R(i) + n_t^P(i) h_t^P(i) \leq 1 \text{ for all } i, t.$$

Moreover, market clearing requires

$$\int_{-\infty}^{+\infty} n_t^s(i) di = N^s \text{ for } s = R, P. \quad (2)$$

The key ingredient of the model is that there is a positive location externality: households like to live in areas where more rich households live. Each location i has an associated neighborhood, given by the interval centered at i of fixed radius γ . Let $H_t(i)$ denote the total space occupied by houses of rich households in the neighborhood around location i ,¹⁸ that is,

$$H_t(i) = \int_{i-\gamma}^{i+\gamma} h_t^R(j) n_t^R(j) dj. \quad (3)$$

Households have non-separable utility in non-durable consumption c and housing services h . The location externality is captured by the fact that households enjoy their consumption more if they live in locations with higher $H_t(i)$. The utility of a household of type s located in location i at time t is given by

$$u^s(c, h, H_t(i)),$$

where $u(\cdot)$ is weakly concave in c and h . For tractability, we assume that u takes the following functional form: $u^s(c, h, H) = c^\alpha h^\beta (A + H)^{\delta^s}$, where α , β , and δ^s are non negative scalars and A is a constant that prevents utility from being zero when H takes the value of zero.¹⁹ Moreover, we assume that $\delta^R \geq \delta^P$, so that rich households who generate the externality benefit from it at least as much as poor households. We want to stress that all of the implications of our model go through even if $\delta^R = \delta^P$.

On the supply side, there is a representative firm who can build housing in any location $i \in (-\infty, +\infty)$. There are two types of houses: rich houses (type R) and poor houses (type P). Each type of household only demands houses of his own type. The marginal cost of building houses of type s is equal to C^s , with $C^R \geq C^P$. If the firm wants to convert houses of type \tilde{s} into houses of type s , he has to pay $C^s - C^{\tilde{s}}$. The (per square foot) price of a house for households of type s in location i at time t is equal to $p_t^s(i)$. Hence there is going to be construction in any empty location i as long as $p_t^s(i) \geq C^s$. Moreover, if the firm wants to construct a house of type s in a location occupied by a house of type \tilde{s} , he has to pay the converting cost and the additional cost of convincing households of type \tilde{s} to leave. Hence, there is going to be construction of houses of type s in any location occupied by agents of type \tilde{s} if $p_t^s(i) \geq C^s - C^{\tilde{s}} + p_t^{\tilde{s}}(i)$.

Finally, there is a continuum of risk-neutral competitive intermediaries who own the houses and rent them to the households. The intermediaries are introduced for tractability. If we allowed the households

¹⁸An alternative is to define the neighborhood externality $H_t(i)$ as the measure of rich households living in the neighborhood around location i (or even as their average income). However, this would make the model less tractable without affecting the substance of the mechanism. A more interesting extension would be to relax the assumption that a neighborhood has a fixed size and make the concept of a neighborhood more continuous. Again the main mechanism of the model would survive this change, but the price schedule would look smoother.

¹⁹Davis and Ortalo-Magné (2010) show that a Cobb-Douglas relationship between housing consumption and non-housing consumption fits the data well along a variety of dimensions.

to own their houses, nothing would change in steady state, but the analysis of a demand shock would be more complicated.²⁰ The (per square foot) rent for a house of type s in location i at time t is denoted by $R_t^s(i)$. As long as the rent in location i at time t is positive, the intermediaries find it optimal to rent all the houses in that location. Also, for simplicity, assume that houses do not depreciate. Competition among intermediaries requires that for each location i the following arbitrage equations hold:

$$p_t^s(i) = R_t^s(i) + \left(\frac{1}{1+r}\right) p_{t+1}^s(i) \text{ for all } t, i, s. \quad (4)$$

3.2 Equilibrium

An equilibrium is a sequence of rent and price schedules $\{R_t^R(i), R_t^P(i), p_t^R(i), p_t^P(i)\}_{i \in R}$ and of allocations $\{n_t^R(i), n_t^P(i), h_t^R(i), h_t^P(i)\}_{i \in R}$ such that households maximize utility, the representative firm maximizes profits, intermediaries maximize profits, and markets clear.

Because of full mobility, the household's maximization problem reduces to a series of static problems. The problem of households of type s at time t is simply

$$\begin{aligned} \max_{c, h, i \in \mathcal{I}_t^s} c^\alpha h^\beta [A + H_t(i)]^{\delta^s}, \\ \text{s.t. } c + h R_t^s(i) \leq y^s, \end{aligned}$$

where households take as given the function $H_t(i)$, the rent schedule $R_t^s(i)$, and the set \mathcal{I}_t^s of locations where houses for type- s households are available. Hence, conditional on choosing to live in location i at time t , the optimal house size is

$$h_t^s(i) = \frac{\beta}{\alpha + \beta} \frac{y^s}{R_t^s(i)} \text{ for all } t, s, i \in \mathcal{I}_t^s. \quad (5)$$

Households choose to live in bigger houses in neighborhoods where the rental price is lower and, conditional on a location, richer households choose bigger houses. Given that households are fully mobile, it must be that at each point in time, the equilibrium rents in different locations make them indifferent. In particular, agents of type s have to be indifferent among living in different locations where houses of their type are available at time t , that is, in all $i \in \mathcal{I}_t^s$.²¹ Then it must be that

$$U_t^s(i) \equiv \alpha^\alpha \beta^\beta \left(\frac{y^s}{\alpha + \beta}\right)^{\alpha + \beta} \frac{[A + H_t(i)]^{\delta^s}}{R_t^s(i)^\beta} = \bar{U}_t^s \text{ for all } t, s, i \in \mathcal{I}_t^s. \quad (6)$$

This, in turn, requires that

$$R_t^s(i) = K^s [A + H_t(i)]^{\frac{\delta^s}{\beta}} \text{ for all } t, s, i \in \mathcal{I}_t^s, \quad (7)$$

for some constant K^s . This expression is intuitive, as rents must be higher in locations with a stronger externality. Moreover, rich households are more willing to pay higher rents for a given locational externality, all else equal.

²⁰When the economy is hit by a positive demand shock, we will show that house prices appreciate by different amounts in different locations. If households own their houses this would introduce an additional source of heterogeneity in wealth which would complicate the analysis.

²¹If there was a location with construction of type s and no type s households living there, the intermediaries would be willing to decrease the rent to 0 inducing households of type s to move into that location.

Proposition 1. *If $\delta^R \geq \delta^P$, there exists an equilibrium with full segregation. If $C^R = C^P$, an equilibrium with full segregation exists if and only if $\delta^R \geq \delta^P$.*

To prove this proposition we proceed by constructing an equilibrium with full segregation, where the rich households are concentrated in the city center, while the poor households live at the periphery of the city. There may be other equilibria with full segregation with more centers of agglomeration of the rich households. It is interesting to notice that, as long as these centers are far enough from each other, the implications in terms of house prices are isomorphic to the equilibrium we focus on.

Let us now proceed to the construction of our equilibrium. As a normalization, we choose point 0 as the center of the city. It follows that $\mathcal{I}_t^R = [-I_t, I_t]$ and $\mathcal{I}_t^P = [-\bar{I}_t, -I_t) \cup (I_t, \bar{I}_t]$, for some $\bar{I}_t > I_t > 0$. Both the size of rich neighborhoods, I_t , and the size of the city, \bar{I}_t , are equilibrium objects. Given that such an equilibrium is symmetric in i , from now on, we can restrict attention to $i \geq 0$.

Since rich households live in locations where there are no poor, it must be that $h_t^R(i) n_t^R(i)$ is either equal to 1 or to 0 and is equal to 1 for all $i \in [0, I_t]$. Then, we can easily derive the externality function $H_t(\cdot)$ as follows:

$$H_t(i) = \begin{cases} 2\gamma & \text{for } i \in [0, I_t - \gamma] \\ \max\{\gamma + I_t - i, 0\} & \text{for } i \in (I_t - \gamma, \bar{I}_t] \end{cases} . \quad (8)$$

That is, neighborhoods close to the city center are richer and enjoy the maximum degree of externality, while the farther a location is from the center the smaller the strength of the externality. Figure 2 shows the externality $H_t(i)$ for a given t as a function of the location. If $\bar{I}_t > I_t + \gamma$, there are going to be locations at the margin of the city where the externality has zero effect. From now on, we assume that the measure of poor households, N^P , is sufficiently large so that $\bar{I}_t > I_t + \gamma$.

Combining (7) and (8), we obtain

$$K^R = R_t^R(I_t) (A + \gamma)^{-\frac{\delta_R}{\beta}} \quad \text{and} \quad K^P = R_t^P(\bar{I}_t) A^{-\frac{\delta_P}{\beta}}, \quad (9)$$

so that we can rewrite the rent schedules as

$$R_t^R(i) = R_t^R(I_t) \left(1 + \frac{\min\{\gamma, I_t - i\}}{A + \gamma} \right)^{\frac{\delta_R}{\beta}} \quad \text{for } i \in [0, I_t], \quad (10)$$

$$R_t^P(i) = R_t^P(\bar{I}_t) \left(1 + \frac{\max\{\gamma + I_t - i, 0\}}{A} \right)^{\frac{\delta_P}{\beta}} \quad \text{for } i \in (I_t, \bar{I}_t]. \quad (11)$$

From the optimizing behavior of the representative firm, it must be that the price of a poor house at the boundary of the city is equal to the marginal cost C^P . Moreover, the price of a rich house at the boundary of the rich neighborhoods must be equal to the price of a poor house, which is the compensation needed to vacate poor households living there, plus the additional cost of transforming a poor house in a rich one. This implies that $p_t^P(\bar{I}_t) = C^P$ and $p_t^R(I_t) = p_t^P(I_t) + C^R - C^P$. In equilibrium prices are constant over time and hence arbitrage conditions (4) require that for each location $i \in \mathcal{I}_t^s$ prices satisfy

$$p_t^s(i) = \frac{1+r}{r} R_t^s(i) \quad \text{for all } t, i, s. \quad (12)$$

Combining these conditions we obtain

$$R_t^P(\bar{I}_t) = \frac{r}{1+r}C^P \text{ and } R_t^R(I_t) = R_t^P(I_t) + \frac{r}{1+r}(C^R - C^P), \quad (13)$$

where, from (7) and (9), we have

$$R_t^P(I_t) = \frac{r}{1+r}C^P \left(\frac{A + \gamma}{A + \max\{\gamma + I_t - \bar{I}_t, 0\}} \right)^{\frac{\delta_P}{\beta}}. \quad (14)$$

Combining these last two expressions with (10), (11), and (12) allows us to determine the rent and the price schedules as a function of I_t and \bar{I}_t only. Figure 2 also shows the shape of the price schedule as a function of the location.

In our full segregation equilibrium, the rich households are concentrated in the city center, while the poor are located at the periphery. Moreover, equilibrium prices reflect the fact that locations that are further away from the center of the rich enclave and closer to the space occupied by poor households are less appealing. In particular, prices are the highest in the center of the rich neighborhoods. As we move away from the center, prices start declining because the space in the neighborhood occupied by rich households goes down. This segregation equilibrium is sustained by the fact that the poor are unwilling to lower their non housing consumption by paying higher rent to get the larger neighborhood externality.

To complete the characterization of the equilibrium, we need to determine the size of the city, \bar{I}_t , and the size of the rich neighborhoods, I_t . Using market clearing (2) together with the optimal housing size (5) and the fact that $\mathcal{I}_t^R = [-I_t, I_t]$ and $\mathcal{I}_t^P = [-\bar{I}_t, -I_t) \cup (I_t, \bar{I}_t]$, we obtain the following expressions for I_t and \bar{I}_t :

$$I_t = \gamma + (A + 2\gamma)^{-\frac{\delta_R}{\beta}} \left\{ \frac{\alpha}{\alpha + \delta} \frac{N^R y^R}{2K^R} - \frac{\beta}{\beta + \delta_R} \left[(A + 2\gamma)^{\frac{\beta + \delta_R}{\beta}} - (A + \gamma)^{\frac{\beta + \delta_R}{\beta}} \right] \right\} \quad (15)$$

$$\bar{I}_t = I_t + \gamma + A^{-\frac{\delta_P}{\beta}} \left\{ \frac{\alpha}{\alpha + \delta} \frac{N^P y^P}{2K^P} - \frac{\beta}{\beta + \delta_P} \left[(A + \gamma)^{\frac{\beta + \delta_P}{\beta}} - A^{\frac{\beta + \delta_P}{\beta}} \right] \right\}. \quad (16)$$

As intuition suggests, the rich neighborhoods cover a larger portion of the city when N^R (the number of rich people) or y^R (the income of rich people) are higher, and when the marginal cost of construction C^R or the interest rate r are lower. Moreover, the city overall is bigger when there are more rich households when the rich households are richer, when there are more poor households or when the poor are richer. Likewise, the city is larger when the marginal cost of construction C^P or the interest rate are lower.

Finally, we have to check that the households choose their location optimally, that is, we have to check that the rich would not prefer to move to a poor neighborhood and vice versa. More precisely, we need to prove that

$$\begin{aligned} U_t^R(i) &\leq \bar{U}_t^R \text{ for all } i \in [I_t, \bar{I}_t] \\ U_t^P(i) &\leq \bar{U}_t^P \text{ for all } i \in [0, I_t] \end{aligned}$$

where $U_t^s(i)$ is defined in expression (7). In the Appendix, we show that both these conditions are satisfied if $\delta^R \geq \delta^P$, completing the proof of the Proposition.

Before proceeding to the analysis of the shocks, let us mention that Proposition 1 claims that an equilibrium with full segregation always exists when $\delta^R \geq \delta^P$, but not that no other equilibria exist. Actually, we can construct another equilibrium with partial segregation, where intervals with only poor people alternate with intervals where poor and rich people coexist. However, we can show that there is no equilibrium with full integration, that is, where poor and rich agents simultaneously live in every occupied location.

3.3 Demand shock

We are now interested in analyzing how house prices, both at an aggregate and at a disaggregate level, react to an unexpected increase in the demand for housing.²² We will do so by focusing on the equilibrium with full segregation that we constructed in the previous section.

In equilibrium, the aggregate price level is given by

$$P_t = \frac{2}{I_t} \int_0^{I_t} p_t^R(i) di + \frac{2}{\bar{I}_t - I_t} \int_{I_t}^{\bar{I}_t} p_t^P(i) di,$$

where, from the analysis in the previous section,

$$p_t^R(i) = \left[C^P \left(1 + \frac{\gamma}{A} \right)^{\frac{\delta_P}{\beta}} + C^R - C^P \right] \left(1 + \frac{\min\{\gamma, I_t - i\}}{A + \gamma} \right)^{\frac{\delta_R}{\beta}} \text{ for } i \in [0, I_t], \quad (17)$$

$$p_t^P(i) = C^P \left(1 + \frac{\max\{\gamma + I_t - i, 0\}}{A} \right)^{\frac{\delta_P}{\beta}} \text{ for } i \in (I_t, \bar{I}_t], \quad (18)$$

with I_t and \bar{I}_t given by (15) and (16).

For concreteness, we analyze the economy's reaction to a migration shock, but the price dynamics are equivalent if we consider any shock that increases housing demand, such as a positive income shock or a reduction in the interest rate. Imagine that at time $t + 1$ the economy is hit by an unexpected and permanent increase in the population N . Let us assume that the measure of both rich and poor households increase proportionally, that is, $N_{t+1}^s = \phi N_t^s$ with $\phi > 1$ for $s = P, R$. We now show that the aggregate level of house prices permanently increases and prices in locations with a higher initial price level typically react less than prices in locations where houses are cheaper to start with and which are closer to the expensive neighborhoods. The new rich households moving into the city want to live close to other rich households, so that the poor neighborhoods close to the rich ones get gentrified and the poor households who used to live there move towards the periphery. This is what we refer to as endogenous gentrification. The house prices in gentrified neighborhoods are driven up due to our externality.

²²The long run reaction of house prices would be symmetric in the case of a negative shock if we introduce some degree of depreciation that is big enough relative to the shock. However, after a negative shock the economy would not jump to the new steady state immediately, but there would be some transitional dynamics. Contact the authors if you are interested in the full analysis of a negative demand shock in the presence of depreciation. Given our data, we only focus on housing booms resulting from positive housing demand shocks.

Let us define the function $g_t(\cdot) : [C^P, \bar{p}] \mapsto [1, \infty)$, where $g_t(p)$ denotes the average gross growth rate between time t and $t + 1$ in locations where the initial price is equal to p , that is,

$$g_t(p) = E_{t+1} \left[\frac{p_{t+1}(i)}{p_t(i)} \mid p_t(i) = p \right].$$

The next proposition shows that after an unexpected permanent positive demand shock, the aggregate price level permanently increases and the price growth rate is higher in locations that had lower price levels initially, whenever prices are higher than the minimum level.²³

Proposition 2. *Imagine that at time $t + 1$ the economy is hit by an unexpected and permanent increase in population, that is, $N_{t+1}^s = \phi N_t^s$ with $\phi > 1$ for $s = P, R$. Then there is a permanent increase in the aggregate price level P_t , and*

$$E_{t+1} \left[\frac{p_{t+1}(i)}{p_t(i)} \mid p_t(i) = \bar{p} \right] < E_{t+1} \left[\frac{p_{t+1}(i)}{p_t(i)} \mid p_t(i) < \bar{p} \right].$$

Moreover, if the shock is large enough, $g_t(p)$ is non-increasing in p for all $p > C^P$.

Figure 3 illustrates the response of house prices in different locations to a positive demand shock (a proportional increase in population). Given that the city is symmetric, the figure represents only the positive portion of the real line. One can notice that both the size of the city \bar{I}_t and the size of the rich neighborhoods I_t expand, and prices remain constant at the two extremes: in the richest locations in the center of the city and far enough away from the rich neighborhoods. Most importantly, prices strictly increase in the rich neighborhoods on the border of the poor neighborhoods where the externality is below its maximum level and, even more, in the poor neighborhoods that are physically close to the rich neighborhoods. Clearly, this makes the aggregate level of prices in the city increase permanently.

The next proposition shows the main implication of our model: among the locations with initial level of price equal to C^P , the ones that appreciate the most are closer to the richer neighborhoods.

Proposition 3. *Imagine that at time $t + 1$ the economy is hit by an unexpected and permanent increase in population, that is, $N_{t+1}^s = \phi N_t^s$ with $\phi > 1$ for $s = P, R$. Then*

$$\frac{d(p_{t+1}(i)/p_t(i))}{di} \leq 0 \text{ for } p_t(i) = C^P.$$

Among the poor neighborhoods, it is the poor neighborhoods in close proximity to the richer neighborhoods that should appreciate the most during a city-wide housing demand shock. This proposition underlies the variation in appreciation rates among the poorer neighborhoods. The poor neighborhoods next to the rich neighborhoods experience large price increases because they gentrify. Rich households expand into the neighborhood thereby increasing the desirability of being in those neighborhoods. This proposition lies at the heart of our following empirical work. An increase in city-wide housing demand -

²³See the online appendix for all the proofs that are not in the text.

perhaps do to an in migration of rich residents - will cause poor neighborhoods on the border of richer neighborhoods to gentrify. The empirical work that follows does, in fact, show strong support for this prediction.

Proposition 4. *Imagine that at time $t + 1$ the economy is hit by an unexpected and permanent increase in population, that is, $N_{t+1}^s = \phi N_t^s$ with $\phi > 1$ for $s = P, R$. Then the growth rate in the aggregate price level is larger the larger is the increase in ϕ and, if the shock is large enough, $d^2 g_t(p) / dp d\phi \geq 0$ for all $p > C^P$ where the derivative is well-defined.*

This proposition shows that if two identical cities are hit by demand shocks of different sizes, the one hit by the larger shock is going to feature both a higher aggregate price growth rate and more price appreciation among the poor neighborhoods due to a higher degree of gentrification. It is also easy to show that two cities with different initial income composition react differently to the same demand shock. In particular, if the shock is large enough, the initially richer city is the one that features both a higher aggregate price growth rate and higher within-city house price convergence.

4 Housing Price Dynamics and Proximity to Rich Neighborhoods: Descriptive Results

In the next few sections, we set out to describe whether within-city house price dynamics can, in part, be explained by the endogenous gentrification mechanism laid out in the previous section. In this section, we further empirically explore the heterogeneity in house price dynamics among poor neighborhoods during a city-wide housing boom and we show that it is the poor neighborhoods that are closest to the rich neighborhoods that appreciate the most.

In order to highlight the importance of the neighborhood consumption externality throughout our empirical work, we need to test it directly against the two other prominent explanations for differential within-city housing price dynamics discussed in the introduction: proximity to jobs and proximity to fixed natural amenities. Both of these classes of models can also generate between neighborhood house price growth differences and neighborhood gentrification in response to city-wide demand shocks. Given our data, it is possible to at least partially distinguish among the different stories. As part of our empirical work, we control directly for variables that measure distance to the center business district (and distance to jobs more broadly) as well as the distance to fixed natural amenities within the city. As we show, it is proximity to rich neighbors irrespective of proximity to jobs and other fixed natural amenities that seem to be driving our results. This is consistent with our modeling choice of focusing on the neighborhood consumption externality.

4.1 Descriptive Specification

To begin, we describe the data by simply estimating the following regression:

$$\frac{\Delta P_{t,t+k}^{i,j}}{P_t^{i,j}} = \mu_j + \beta_1 \ln(\text{Dist}_t^{i,j}) + \Gamma X_t^{i,j} + \Psi Z_t^{i,j} + \epsilon_{t,t+k}^{i,j} \quad (19)$$

where $\Delta P_{t,t+k}^{i,j}/P_t^{i,j}$ is the growth in housing prices between period t and $t+k$ within neighborhood i in city j . Across our various specifications, we use either the Case-Shiller, Zillow or Census data to compute neighborhood housing price growth. When using the Case-Shiller and Zillow data, our measure of a neighborhood is a zip code. When using the Census data, our measure of a neighborhood is either a zip code or a census tract. Given that neighborhoods within a city have different amounts of homeowners or potential housing market transactions, all regressions are weighted by the number of owner-occupied housing units in the neighborhood during period t . Finally, we wish to stress that all of our regressions also include city fixed effects, μ_j . As a result, all of our identification comes from within-city variation. We report heteroscedasticity robust standard errors that are clustered at the city level.

When estimating the above regression, our sample only includes low housing price neighborhoods within the city. We define low housing price neighborhoods as those neighborhoods whose median housing price at time t is in the bottom half of neighborhoods with respect to median housing prices across all neighborhoods in city j at time t .²⁴ The reason for this restriction is that our theory predicts that it is the poorer neighborhoods that are close to richer neighborhoods that should appreciate more than other poorer neighborhoods in response to city-wide housing demand shocks as richer residents spill over into these neighborhoods. As in Section 2, we use the Census data to define the level of period t median housing prices for each neighborhood when segmenting the sample.

The variable of interest in the above regression is β_1 , the coefficient on $\ln(\text{Dist}_t^{i,j})$, where $\ln(\text{Dist}_t^{i,j})$ measures the log of the distance (in miles) to the nearest zip code in the city that resides in the top quartile of neighborhoods with respect to median housing prices in period t .²⁵ In essence, we are asking whether a poor neighborhood which is next to a high price neighborhood appreciates more than an otherwise similar poor neighborhood farther from the high price neighborhood.

The vector $X_t^{i,j}$ includes a series of variables designed to control for initial differences across the neighborhoods. These controls include the log of median household income of residents in neighborhood i in period t , the log of the median initial house price of residents in neighborhood i in period t , the fraction of the residents in neighborhood i in period t that are African American, and the fraction of the residents in neighborhood i in period t that are Hispanic. When the Census data is used to compute housing price appreciations, we also include a vector of variables to proxy for the change in structure quality within the

²⁴Sometimes in the text we will refer to these neighborhoods as “poor neighborhoods”. We do this for expositional ease. We also used an income based measure to define poor neighborhoods. Given the very high correlation between neighborhood average income and neighborhood housing prices, the results are broadly consistent if we segment neighborhoods by initial income as opposed to initial house prices.

²⁵We measure distance from the centroid of each neighborhood.

neighborhood between t and $t+k$. These controls include: the change in the fraction of homes in the tract that are single-family-detached, the change in the fraction that have zero or one bedrooms, the change in the fraction that have two bedrooms, the change in the fraction that have three bedrooms, the change in the fraction built in the past 5 years, the change in the fraction built between 5 and 20 years ago, the change in the fraction built between 20 and 40 years ago, and the change in the fraction built between 40 and 50 years ago.

We also include a vector $Z_t^{i,j}$ which is designed to control for the other potential mechanisms which can generate differential price movements across neighborhoods within a city. Specifically, we control for the average distance to the closest center business district within the city as reported by the 1982 Census of Retail Trade.²⁶ Additionally, the Census data provides an additional measure of proximity to jobs in that it tracks how long it takes for individuals in the neighborhood to get to work. Given this, we also include the median commuting time of individuals within neighborhood i during period t as an additional control. Finally, we control for the distance to fixed natural amenities like lakes, rivers, and oceans that either are part of the city or border the city.

4.2 Descriptive Results

Table 3 shows the results of the above regressions using different time periods, different housing price appreciation measures, and different levels of aggregation. The first three columns show the results for the 2000-2006 period where we use Case-Shiller (columns 1 and 2) and Zillow (column 3) data. In column 1, we exclude the Z vector of controls to gauge their impact while in columns 2 and 3, we include both the X and Z vectors of controls. We run the specifications shown in columns 1 - 3 on the same sample: all zip codes within Case-Shiller cities for which both a Case-Shiller index and a Zillow index exists and which were in the bottom half median house prices within the city in 2000. There are 211 such zip codes.

As shown above, there is much heterogeneity with respect to housing price appreciation among poor neighborhoods within a city when the city experiences a city-wide housing price boom. The results in columns 1 - 3 show that there is some systematic variation among the poor neighborhoods. In particular, it is the initially low price neighborhoods in 2000 which were in close proximity to the high price neighborhoods that appreciated at a much higher rate than otherwise similar initially low price neighborhoods that were farther away. These results hold even after controlling for proximity to the city's Central Business District (CBD), average commuting times, distance to fixed natural amenities and the X vector of neighborhood controls (column 2).

In terms of economic magnitudes, the estimates are non-trivial. For example, the results in column 2 suggest that low priced neighborhoods that were roughly 4 miles away from higher price neighborhoods appreciated at 12.8 percentage point lower rates than low priced neighborhoods that were roughly 1 mile

²⁶The CBD data can be found at <http://www.census.gov/geo/www/cbd.html>.

away from higher priced neighborhoods ($0.064 * 2$, p-value < 0.01). Given that the average neighborhood house price appreciation rate for the neighborhoods in our sample during this period was roughly 80 percent, the estimated relationship with distance to high price neighborhoods is non trivial.

Two other results are worth noting from columns 1 - 3 of Table 3. First, adding in controls for proximity to CBD, mean commuting time, and proximity to lakes, rivers, and oceans within the city has very little effect on the estimated coefficient (column 1 vs. 2). The results shown are for the Case-Shiller data, but also hold for the Zillow data. To the extent that the coefficient changes, it is driven by proximity to the CBD. Conditional on controlling for proximity to the CBD, proximity to lake front, river front, or ocean front has little, if any, affect on the estimated coefficient. While not conclusive, it is comforting that our estimated coefficients do not change when we include proximity to the fixed natural amenities that we do observe suggesting that the ones we do not observe may not have much impact on our estimates. Second, the estimated coefficients are similar in magnitude regardless of whether we use the Case-Shiller data or the Zillow data (column 2 vs. 3). Again, it is reassuring that our descriptive findings show up when we use both the Case-Shiller data and the Zillow data, given that each has its own strengths and weaknesses.

In columns 4-7, we show similar results for the 1990-2000 period. All specifications in these columns control for both the full vector of X and Z controls. In columns 4 and 5, we use the Case-Shiller data on Case-Shiller zip codes. The difference between the two columns is that in column 5 we also include an additional regressor: $\ln(Dist_t^{i,j}) * Bust_{t,t+k}^j$ where $Bust_{t,t+k}^j$ is an indicator variable taking the value of one if city j experienced non-positive housing price growth between t and $t + k$. The reason for the inclusion of this variable is that our theory says that for cities that experience a positive housing demand shock, the neighborhoods bordering the rich neighborhoods should appreciate more than the neighborhoods far from the rich neighborhoods. We do not include this variable in the 2000-2006 period because all cities in our sample experienced a positive house price increase. However, as seen from Appendix Table A3, some cities in our Case-Shiller sample experienced real housing price declines during the 1990s. The relationship between house price growth among poor neighborhoods close to and far from high price neighborhoods should differ depending on whether the city experienced a positive or non-positive housing demand shock during the period. As seen by comparing columns 4 and 5 of Table 3, the cities which experienced a non positive housing price boom had very little difference in housing price growth between poor neighborhoods that were close to high price neighborhoods and poor neighborhoods that were farther away ($-0.071 + 0.085$). However, for cities that experienced a positive city-wide housing price increase, the estimated relationship in the 1990s mirrored what we found in the 2000 - 2006 period. The estimated coefficient on log distance during the 1990s was -0.071 (p-value = 0.01).

In columns 6 and 7, we show that the results are roughly consistent using the Census data during the 1990-2000 period. We define neighborhoods as census tracts and use different samples. In column 6, the sample is all census tracts within only the 24 Case-Shiller cities. In column 7, the sample is census tracts

in any U.S. city that has at least 30 census tracts contained within the city. There were 172 U.S. cities in 1990 that met this condition.²⁷ Column 7 shows that the main results still hold when examining price movements at the level of census tracts and that the results are not simply limited to Case-Shiller cities.

Finally, in column 8, we show the results for the 1980-1990 period. The specification in column 8 is analogous to the one shown in column 7 aside from the fact that it looks at the 1980-1990 period for the 109 U.S. cities in 1980 that had at least 30 census tracts contained within the city. The patterns in the 1980s are similar to those found in the 1990s and early 2000s. It is important that the results are similar between the Case-Shiller, Zillow, and Census housing price measures. In the next section, we explore the response of within-city house price dynamics to exogenous city-wide housing demand shocks. To get enough power, we need to use the large samples shown in columns 7 and 8 for which only the Census housing price measures are available. The price responses are slightly higher for the census tract specifications (columns 6 - 8) relative to the zip code specifications (columns 1 - 5).

Before moving to the next section, we wish to emphasize a few results from Table 3. First, we view these results as offering a new descriptive set of facts about within-city house price movements. Not only do poor neighborhoods, on average, appreciate more than higher price neighborhoods (documented above), it is the poor neighborhoods in close proximity to the higher priced neighborhoods that appreciate the most within the city when the city as a whole is experiencing positive housing price growth. These results hold controlling for distance to the CBD, commuting times, and distance to fixed natural amenities. As far as we know, this fact is also new to the literature. Second, this fact is extremely robust across measures of house price growth, across different levels of neighborhood aggregation, and across different time periods spanning three decades.²⁸ Finally, the results are consistent with the theory we outlined above. If households get utility from proximity to rich neighbors, the spatial equilibrium model outlined above predicts that city level shocks which cause housing demand to increase for rich households will result in poorer neighborhoods on the border of rich neighborhoods appreciating to a much greater extent than otherwise similar poorer neighborhoods farther away.

5 Housing Price Dynamics and Proximity to Rich Neighborhoods: Exogenous Housing Demand Shock

In the previous section, we established empirical relationships that are consistent with our theory of endogenous gentrification. However, with the prior descriptive results we cannot make any claim about causation, given that we are silent on the shocks behind the housing demand change. In this section,

²⁷In the online robustness appendix, we specify in detail all our sample criteria when using the expanded set of census tracts. In particular, we discuss how we select census tracts that are consistently defined over time.

²⁸We also performed a series of additional robustness specifications on our results. These results are shown in our online robustness appendix. One such robustness specification is to include controls for the age of the housing stock within the neighborhood during period t . As shown by Brueckner and Rosenthal (2008), the age and quality of the housing stock could be an important determinant of which neighborhoods will subsequently gentrify. Our results are robust to the inclusion of the initial age of the housing stock in our specifications.

we directly explore whether an exogenous shock to housing demand in a city affects property prices differentially across neighborhoods within that city in a way that is consistent with our theory. As before, we also rule out that these effects are being driven by proximity to jobs or proximity to fixed natural amenities.

5.1 Exogenous Housing Demand Shock: Bartik Instrument

To measure exogenous shocks to local housing demand for each city j between t and $t + k$, we use the variation in national earnings by industry between t and $t + k$. This approach of imputing exogenous income shocks for local economies was developed by Bartik (1991) and has been used extensively by others in the literature as a measure of local labor demand shocks.²⁹ In doing so, we are explicitly equating local income shocks with local housing demand shocks. As shown by Blanchard and Katz (1992), such positive Bartik-type local income shocks causes an influx of population from other cities which puts upward pressure on city-wide housing prices.

To compute exogenous changes in city-level income, we use the initial industry composition of residents within the city's MSA in period t . Note, that even though we are examining house price dynamics within a given city, our estimate of the local demand shock is based on the industry mix of the MSA as a whole. Given the amount of commuting into and out of the city from the suburbs (in both directions), we feel the MSA income shock is a broader proxy for city-wide changes in housing demand. Then, for each MSA ^{j} , we compute the predicted income growth for the MSA using the initial industry shares and the growth in income for individuals in those industries between t and $t + k$ for the entire U.S. (excluding residents from MSA ^{j}).

For our results examining the neighborhood response to local housing demand shocks, we focus our attention on the 1980-1990 period. We analyze this period because there is significant variation across MSAs in predicted MSA-wide average income growth based on industry composition during this time. Specifically, we use the five percent samples from the 1980 and 1990 IPUMS data to compute MSA-level predicted income growth. For the procedure, we use two-digit industry classifications. Our measure of income is individual earnings. The only restrictions we place on the data are that the individual had to be employed full-time (worked 48 weeks or more in the prior year and usually worked more than 30 hours per week) and had to be between the ages of 25 and 55. Again, when computing the predicted income growth for a given MSA, we exclude the residents of that MSA when computing the national income growth between 1980 and 1990 for each of the industries. To compute the MSA predicted income growth, we simply multiply the industry growth rate in earnings by the fraction of people between 25 and 55 in each MSA working full time in those industries in 1980.

There is a large variation in actual income growth by industry between 1980 and 1990. For example,

²⁹See, for example, Blanchard and Katz (1992), Notowidigdo (2010), and Diamond (2010).

the Security, Commodity Brokerage, and Investment Companies industry had a real appreciation of annual earnings of roughly 59% during the 1980s. Likewise, the Legal Services industry had a real appreciation of annual earnings at 55%. On the other hand, the Trucking Services and Warehousing and Storage industry only had a real appreciation of annual earnings of 3%. As a result of differences in industry mixes across MSAs, there is a nontrivial amount of predicted income variation across the MSAs. For our results in this section, we include all MSAs that contain a city which has at least 30 census tracts within the city. This will be the same sample that we used in column 8 of Table 3. For those MSAs, the mean predicted income growth during the 1980s using the Bartik measure was 19.4 percent with a standard deviation across the MSAs of 1.5 percent.

As shown by others in the literature, the predicted Bartik income growth for the MSA does in fact predict actual MSA income growth. A simple regression of actual MSA income growth during the 1980s on predicted MSA income growth during the 1980s yields a coefficient on predicted MSA income growth of 1.95 (with a standard error of 0.58) and a F-stat of 11.13 .

5.2 Specification

To examine the effect of exogenous housing demand shocks on within-city house price dynamics, we estimate a specification similar to (19):

$$\frac{\Delta P_{t,t+k}^{i,j}}{P_t^{i,j}} = \mu_j + \beta_1 \ln(Dist_t^{i,j}) + \beta_2 \ln(Dist_t^{i,j}) * \widehat{IncShock}_{t,t+k}^j + \Gamma X_t^{i,j} + \Psi Z_t^{i,j} + \epsilon_{t,t+k}^{i,j} \quad (20)$$

where $\Delta P_{t,t+k}^{i,j}/P_t^{i,j}$, $\ln(Dist_t^{i,j})$, $Z_t^{i,j}$, $X_t^{i,j}$ and μ_j are defined as above. The variable $\widehat{IncShock}_{t,t+k}^j$ denotes the predicted income growth for the MSA corresponding to city j between t (1980) and $t+k$ (1990) based on the industry mix of residents in the MSA in 1980. We are interested in β_2 , the coefficient on the interaction term. With this regression, we are asking whether, for a given sized city income shock, poor neighborhoods within the city in close proximity to rich neighborhoods within the city appreciate more than otherwise similar poor neighborhoods that are farther away. For our measure of housing price growth, we use the Census data and for the measure of neighborhood we use census tracts. As with the results above that use census house price data, we include controls for changes in the neighborhood housing stock characteristics as part of our X vector. Otherwise, the X vector is the same. Because we are including an estimated regressor in our estimation of (20), we bootstrap our standard errors.³⁰

5.3 Results

The results of estimating the above equation are shown in Table 4. In column 1, we estimate (20) as it is specified. The variable of interest is in the second row and provides an estimate of β_2 . As with the

³⁰For each iteration in the bootstrap procedure, we sample with replacement from the 1980 and 1990 IPUMS data to calculate the first stage. Next, we sample with replacement from the census tracts in each of the cities in our sample to calculate the second stage. We report standard errors calculated by repeating this process 2,500 times. We also estimated the standard errors using only 500 repetitions. The results were very similar.

simple descriptive results shown in Table 3, an exogenous shock to city income results in house prices increasing more in poor neighborhoods that are in close proximity to rich neighborhoods (coefficient = -2.27 with a standard error of 0.53). To help interpret the economic magnitude, we consider the differential housing price response to otherwise similar poor neighborhoods which are 1 and 4 miles away from a rich neighborhood in response to the MSA receiving a one-standard deviation MSA-level Bartik shock. Given the estimated coefficient, a census tract that starts in the bottom half of the city-wide house price distribution in 1980 appreciates by 6.8 percent more when they are 1 mile from a high price census tract relative to an otherwise similar census tract that is 4 miles away ($-2.27 * 2 * 0.015$). This result is non-trivial and is in line with the general descriptive patterns shown in Table 3.

In column 2, we re-run the same specification replacing the log distance variable with two dummy variables measuring the proximity to high housing price census tracts. We do this to explore in greater depth whether the relationship between housing price growth and proximity to richer neighborhoods declines monotonically as the poor neighborhoods become farther away from the rich neighborhoods. Specifically, we replace the log distance measure with dummies indicating whether the census tract was between 0 and 1 miles and between 1 and 3 miles, respectively, to the nearest census tract in the top quartile of the city-wide house price distribution in 1980. We are interested in the coefficient on the interaction between these dummies and the predicted income shock. The house price response to an exogenous city-wide income shock is positive and statistically different from zero for both distance ranges. Reassuringly, the house price response is four times as large for census tracts that are between 0 and 1 miles from the high housing price neighborhoods relative to census tracts between 1 and 3 miles from the high housing price neighborhoods (p-value of difference of the two coefficients = 0.02). Given the average size of census tracts, almost all the initially poor census tracts within 1 mile of a rich census tract actually abut the rich neighborhood. In other words, the biggest responses in prices within a city to a city-wide housing demand shock are for those poor census tracts that border richer census tracts. The estimated magnitudes are also nontrivial. In response to a one standard deviation Bartik shock, poor census tracts within 0 and 1 miles and within 1 and 3 miles appreciated at 6.1 and 1.5 percent more, respectively, than poor census tracts more than 3 miles away.

We wish to make four additional comments about the results in Table 4. First, given that we are including city fixed effects, all our results are identified off of within-city variation. Second, as with the results in Table 3, we are controlling for proximity to CBD, average commuting times, and proximity to fixed natural amenities. Given this, our results are being driven by proximity to rich neighborhoods above and beyond proximity to the center business district or fixed natural amenities within the city. Third, although not shown, the results hold broadly for the 1990s as well but power is more of an issue during that time period. Finally, we explored whether the responsiveness of house prices in poor neighborhoods that were close to rich neighborhoods to Bartik shock was greater in cities where housing supply was more

inelastic. To do this, we further interacted our distance to rich neighborhoods and our distance measure multiplied by the Bartik shock with Saiz’s measure of MSA housing supply elasticity (Saiz, 2010). The point estimates of the triple interaction indicated that the price response of poor neighborhoods bordering rich neighborhoods was stronger in more inelastic cities. However, the standard error was much too large to be conclusive.

6 Housing Price Dynamics, Proximity to Rich Neighborhoods, and Evidence of Neighborhood Gentrification

The results in Tables 3 and 4 show that during city-wide housing price booms it is the poor neighborhoods that are in close proximity to rich neighborhoods that appreciate the most. In this section, we examine more deeply the mechanism of our model. Our model predicts that the poor neighborhoods next to the rich neighborhoods are the ones that appreciate the most because they are the ones where rich households migrate to after a city-wide housing demand shock. This implies that neighborhoods that experience higher house price appreciation should also show signs of gentrification.

This section proceeds in three parts. First, we show a simple relationship between housing price growth within the neighborhood between t and $t+k$ and income growth within the neighborhood during the same time period. Second, we explore the spatial patterns of neighborhoods that, ex-post, have gentrified. In particular, we show that poor neighborhoods that are observed to have gentrified are much more likely to border richer neighborhoods than otherwise similar poor neighborhoods. Finally, and most importantly, we use our Bartik-style instrument to show that when a city experiences an exogenous income shock, the poor neighborhoods on the border of rich neighborhoods are much more likely to gentrify than poor neighborhoods that are farther away from rich neighborhoods.

6.1 Descriptive Relationship of Neighborhood House Price Growth and Neighborhood Income Growth

To analyze whether neighborhoods that experienced a rapid growth in prices also experienced signs of gentrification, we estimate the following simple relationship:

$$\frac{\Delta P_{t,t+k}^{i,j}}{P_t^{i,j}} = \mu_j + \beta \frac{\Delta Y_{t,t+k}^{i,j}}{Y_t^{i,j}} + \epsilon_{t,t+k}^{i,j}$$

where $\Delta Y_{t,t+k}^{i,j}/Y_t^{i,j}$ is income growth in neighborhood i of city j during the period t to $t+k$ and where $\Delta P_{t,t+k}^{i,j}/P_t^{i,j}$ and μ_j are defined as above. The regression asks whether or not a neighborhood that experiences higher house price growth than other neighborhoods within the city also experiences higher income growth than other neighborhoods within the city. In this specification, we are equating neighborhood income growth with neighborhood gentrification. We know that neighborhood gentrification is usually associated with poor neighborhoods experiencing an increase in resident income. In the work below, we try to explore different and perhaps broader measures of neighborhood gentrification. However, given

that the force in our model that drives the house price appreciation of poor neighborhoods that abut rich neighborhoods in response to city-wide housing demand shocks is the exodus of poor residents which are replaced by richer residents, we think neighborhood income growth is a good summary statistic for the mechanism we are trying to highlight.

The results of this regression, using different samples, different measures of housing price growth and different levels of aggregation for a neighborhood, are shown in Table 5. In columns 1- 4, we look at the relationship between house price growth and income growth across neighborhoods during the 1990s. In column 5, we explore the relationship during the 1980s. In columns 1 and 2, we restrict our analysis to Case-Shiller zip codes using Case-Shiller house price data (column 1) and Census house price data (column 2), respectively, to compute housing price growth. In columns 3 - 5, we only use Census house price data to compute housing price growth and define neighborhoods at the level of a census tract. In column 3, we explore census tracts in Case-Shiller cities. In columns 4 and 5, we explore all census tracts in cities that have at least 30 consistently measured census tracts over the decade. As with our analysis in Tables 3 and 4, we also restrict all samples to include only neighborhoods within the city that are in the bottom half of the city's house price distribution in period t . As a result, the samples used for the columns in Table 5 are analogous to samples shown in columns 4-8 of Table 5.

The main take away from Table 5 is that there is a strong relationship between neighborhood income growth and neighborhood house price growth regardless of the house price measure, regardless of the level of aggregation and regardless of the time period. Although not shown, these results hold if we include the X vector of initial neighborhood controls into the regression. Not surprisingly, neighborhood income growth is strongly positively correlated with neighborhood house price growth.

6.2 Gentrification and Proximity to High Income Neighborhoods An Ex-Post Analysis

In the prior subsection, we looked at the relationship between income growth and house price growth across neighborhoods within a city. The poor neighborhoods that appreciated the most within the city during a given time period also experienced, on average, large increases in neighborhood income during the same time period. In this subsection, we perform a different analysis to highlight the spatial nature of gentrification. In particular, we identify all neighborhoods within cities that ex-post can be classified as having gentrified according to some broad definition and we examine their spatial proximity to high income neighborhoods. Our goal is to illustrate that when gentrification occurs, it almost always occurs in poorer neighborhoods that border higher income neighborhoods.

For our analysis, we define a gentrifying neighborhood as a neighborhood within a city that starts with median neighborhood house prices in the bottom half of the city's house price distribution in period t and where the median real income of neighborhood residents grows by either 50 percent or 25 percent between t and $t + k$. We use our broadest sample of cities with at least 30 consistently measured census

tracts for the 1980s and 1990s. Specifically, the samples we use are the same as those used in columns 7 and 8 of Table 3.

For this analysis, we simply regress a dummy variable for whether the initially poor neighborhood gentrified by some income growth metric on distance dummies to the nearest rich neighborhood and city fixed effects. As above, we define rich neighborhoods as those neighborhoods that were in the top quartile of the city's median house price distribution in period t . We define four dummy variables to measure the poor neighborhood's proximity to the rich neighborhoods: between 0 and 0.5 miles, between 0.5 and 1 mile, between 1 and 2 miles, and between 2 and 3 miles. Finally, we run such regressions separately for two measures of gentrification: neighborhood real income growth greater than 50 percent during the decade and neighborhood real income growth greater than 25 percent during the decade.

The results of these regressions are shown in Table 6. The results are quite striking. Take, for example, the results where gentrification is defined as a poor neighborhood having average neighborhood income growth increasing by at least 50 percent during the decade. Between 1980 and 1990, 11% of all poor neighborhoods gentrified by this metric. The comparable number between 1990 and 2000 was 5.9%. During the 1980s, the probability of gentrification was 6.9 percentage points higher if the census tract was between 0 and 0.5 miles from a high house price neighborhood than for a poor census tract that was more than 3 miles away from a rich neighborhood (column 1 of Table 3, p -value < 0.01). The coefficient is large in economic magnitude. Given the base gentrifying rate was 11 percentage points, a poor census tract being within 0 and 0.5 miles was associated with a 63 percent increase in the probability of gentrification. During the 1990s, poor census tracts that were within 0.5 mile of a rich census tract were 96.6 percent more likely to gentrify than poor census tracts that were more than 3 miles away from the rich census tracts. Similar patterns are found in both decades if we define gentrification as neighborhood income growth increasing by 25%.

Poor census tracts that were within 0.5 miles of a rich census tract almost always abutted the rich census tract. Among the poor census tracts, as one moves farther away from the nearest rich census tract, the probability of gentrification declines monotonically. These results are also seen in Table 6.

The results in this subsection are consistent with the housing price dynamics in our model. Poor neighborhoods tend to gentrify only when they are in close proximity to existing rich neighborhoods. The results show that there is definitely a spatial nature to the gentrification process. When poor neighborhoods experience large income growth, these neighborhoods almost always border an existing rich neighborhood.

6.3 Gentrification and Proximity to High Income Neighborhoods: Exogenous Housing Demand Shock

In the final subsection, we complete our analysis by assessing whether exogenous city-wide demand shocks cause poor neighborhoods in close proximity to richer neighborhoods to endogenously gentrify. As we saw in Table 4, poor neighborhoods in close proximity to rich neighborhoods had house price growth that was

larger than other poor neighborhoods in response to positive city-wide Bartik shocks. In this subsection, we show that these close neighborhoods were also more likely to experience signs of gentrification.

To look for signs of endogenous gentrification, we estimate the following:

$$G_{t,t+k}^{i,j} = \mu_j + \beta_1 \ln(\text{Dist}_t^{i,j}) + \beta_2 \ln(\text{Dist}_t^{i,j}) * \widehat{\text{IncShock}}_{t,t+k}^j + \Gamma X_t^{i,j} + \epsilon_{t,t+k}^{i,j} \quad (21)$$

where $G_{t,t+k}^{i,j}$ is a measure of gentrification in neighborhood i of city j between t and $t+k$. Specifically, we use three measures of $G_{t,t+k}^{i,j}$: the percentage increase in neighborhood income between t and $t+k$, the percentage point change in the poverty rate between t and $t+k$, and the percentage point change in the fraction of residents in the neighborhood who had a college degree or more. Aside from the change in the dependent variable, (21) is analogous to (20) estimated above for neighborhood housing price growth. Moreover, the sample and definition of $\widehat{\text{IncShock}}_{t,t+k}^j$ are exactly the same as the specifications used to estimate (19) in Table 3.

The results from estimating (21) are shown in Table 7. In response to a city-wide housing demand shock, it is the poor census tracts that are in close proximity to the rich census tracts that are much more likely to experience rising incomes, declines in the poverty rate, and rising educational attainment of residents relative to poor census tracts that are farther from the rich census tracts. Specifically, in response to a one-standard deviation Bartik shock, poor census tracts that were 1 mile from rich neighborhoods experienced income growth that was 1.7 percentage points higher than poor neighborhoods that were 4 miles away. Given that the average census tract in our sample experienced income growth of 14.9 percent during the decade, this represents an increase in income of 11.4 percent for poor neighborhoods that are close to rich neighborhoods in response to a one standard deviation Bartik income shock. Likewise, poor neighborhoods that are 1 mile from the rich neighborhoods experienced 23 percent lower increases in the poverty rate and 25 percent higher increases in the fraction of residents with a college degree or more relative to otherwise similar poor neighborhoods that are 4 miles from the rich neighborhoods.

7 Robustness and Additional Discussion

In this section, we discuss a few outstanding issues with our model and associated empirical work described above. Also, we discuss how our results complement and extend many existing literatures. Finally, we discuss how our within-city results can offer insights into cross-city housing price dynamics.

7.1 Model and Empirical Extensions

7.1.1 Rents vs. House Prices

The main implication of our model is that the gentrification process should increase equilibrium rents within the gentrifying neighborhood. Yet, all of our empirical tests explore the response of changing house prices. House price dynamics can diverge from rental price dynamics if there are changes to the

rate at which rents are discounted. To completely rule out that changes in discount rates are driving our empirical work, one would want to explore the response of rental prices across neighborhoods within the city in response to a city-wide housing demand shock. Systematic constant quality rental price indices, however, are not available at the neighborhood level for most cities. Given that, it is hard to explicitly test rental price dynamics at the neighborhood level in response to city-wide housing demand shocks.

Despite the lack of systematic data, we do not believe that changes in discount rates (as opposed to changes in rental prices) are driving our results. The first reason for this is that Morris Davis - our discussant on the paper at the NBER Winter 2010 EFG meeting - documented that many of our broad empirical patterns hold using rental price data that he purchased for the Chicago metro area. He only purchased the data for the Chicago metro area and the level of neighborhood aggregation was much larger than a zip code. These results suggest that, at least for neighborhoods in Chicago, variation similar to our model predictions is showing up as variation in rents.³¹

Second, and potentially more importantly, most of our main empirical tests focus on two otherwise similar poor neighborhoods within a city. Aside from the proximity to rich neighborhoods, the neighborhoods we compare are similar on observables in the initial period. It is hard to think of a reason why the discount rate between two otherwise similar neighborhoods within a city would differ aside from the expectation of future rental price growth that we are emphasizing within our model.

Some readers of our paper have speculated that the extension of credit to low income borrowers could generate the results we document. In order for that to be the case, the story needs to be more complicated than just the extension of credit to low income borrowers. All of our identification comes from comparing neighborhoods with poor residents that are close to rich neighborhoods and neighborhoods with equally poor residents that are farther away. The extension of credit to low income borrowers which occurred recently should have equally affected poor residents in neighborhoods close to rich neighborhoods and poor residents in neighborhoods farther away. Additionally, the patterns we document hold in periods when aggregate discount rates were changing a lot (during the 2000s) and periods when aggregate discount rates were relatively constant (during the 1990s). Given that our results are robust in all time periods we analyze, it becomes harder to tell a story where changes in aggregate discount is the common driver of our results. Further, most of our main Bartik specifications were estimated off of data from the 1980s and 1990s. The extension of credit to low income borrowers was not as pronounced during these periods as it was during the 2000s.

We definitely agree that if systematic high quality within-city rental price data becomes available, it would be beneficial to redo our main empirical specifications with rental prices as opposed to housing prices. However, given that many of the patterns we document are found using rental prices within Chicago and given that we primarily estimate our results within a city holding initial neighborhood characteristics

³¹Morris's discussion of our paper can be found on his web-page: <http://morris.marginalq.com/#Research>.

constant and that we find similar results across all time periods that we analyze, we do not believe that changes in discount rates are the main factor driving our results.

7.1.2 Expected vs. Unexpected Housing Demand Shocks and Uncertainty

In our theoretical analysis, we have assumed that housing demand shocks were unexpected and permanent and that there was no uncertainty. One question that naturally arises is how would our model differ if the housing demand shocks were ex-ante expected, if the housing demand shocks were temporary, or if there was some uncertainty surrounding the housing demand shocks.

In our model, expectations of future demand shocks would be capitalized into current house prices. Finding out today that a neighborhood will gentrify tomorrow in response to a known city-wide housing demand shock will result in an increase in house prices in that neighborhood today even if the underlying composition of the neighborhood remained constant. At the time of the actual city-wide housing demand shock, there would be no subsequent change in house prices in the gentrifying neighborhood because the expectation of gentrification was already capitalized into house prices. As a result, to the extent that housing demand shocks are expected, our estimates in Table 4 of neighborhood house prices to contemporaneous city-wide housing demand shocks would be downward biased relative to the true effects. Given that the Bartik shocks could have been expected, we would be underestimating the true relationship between city-wide demand shocks and neighborhood house price growth. The fact that we are finding some results in our IV specifications suggests that at least a portion of the Bartik housing demand shock was unexpected. This is consistent with the MSA-wide housing price responses to the Bartik shocks shown in Blanchard and Katz (1992).

How predictable are the estimated Bartik shocks? We can address this question by regressing our predicted MSA Bartik shocks between 1980 and 1990 on MSA characteristics in 1980 (including 1980 MSA income) and actual MSA income growth between 1970 and 1980. We run this specification on the roughly 90 MSAs from which our census tracts are drawn in the sample used in column 8 of Table 3. The R-squared of the regression of our MSA-level Bartik shock on 1980 actual MSA-level average family income and the growth rate in actual MSA-level average family income between 1970 and 1980 was only 0.015. Neither of the regressors were statistically significant. From this, we conclude that there is, at best, very little predictability to the Bartik shocks we use to isolate city level housing demand shocks.

It would be interesting to introduce uncertainty in our model to study whether changes in expectations of city-wide housing demand can generate what look like neighborhood housing bubbles. For example, if one believes the city is going to be hit by a positive housing demand shock, the neighborhoods that are most likely candidates for gentrification should experience a jump in housing prices. The housing prices in these neighborhoods would increase even though there was no change in neighborhood characteristics. If the expectation of the demand shock ended up being overly optimistic, prices in these neighborhoods

would revert downward. Our gentrification model with expectations would then yield predictions about within-city variation in housing price booms and busts to changes in expectations about city-wide housing demand shocks. Mian and Sufi (2009) have found that sub prime neighborhoods within a city experienced rapid house price growth relative to other neighborhoods within the city. They attribute the patterns in the data to the extension of credit to sub prime borrowers. Further, they find that actual income did not increase in these neighborhoods allowing them to rule out that these neighborhoods experienced increases in house prices because of actual neighborhood gentrification. However, the expectation of future gentrification could also explain the patterns they document. We think that extending our model to include expectations of future housing demand shocks and then testing the empirical predictions of the model would be an interesting area for future work.

7.2 Related Literature

Our work complements and extends existing literatures dealing with neighborhood gentrification, the importance of neighborhood consumption externalities in determining house prices, and within-city house price dynamics. In this sub-section, we briefly highlight the relationship of our work to these other literatures. Again, the main goal of our paper is to show that city-wide housing demand shocks cause endogenous gentrification and to link this phenomenon to neighborhood and city-wide house price dynamics. We view this as a novel and important contribution to these existing literatures.

7.2.1 Neighborhood Gentrification

There is a large and dispersed literature on neighborhood gentrification. For a recent excellent review, see Kolko (2007). Some of this literature emphasizes correlates with neighborhood gentrification. For example, both Kolko (2007) and Brueckner and Rosenthal (2008) emphasize that the age and quality of the housing stock within a poor neighborhood is an important predictor of whether or not that poor neighborhood ever gentrifies. In a separate strand of work, Kolko (2009) shows that proximity to jobs is also correlated with the propensity for a neighborhood to gentrify. Given these well documented results, we show above that our main empirical results are robust to the inclusion of the initial age and quality of the neighborhood housing stock as well as proximity to jobs as additional regressors.

Additionally, there is some work that emphasizes the importance of spatial dependence - either theoretically or empirically - in predicting neighborhood gentrification. For example, Brueckner (1977) finds that urban neighborhoods in the 1960s that were in close proximity to rich neighborhoods got relatively *poorer* between 1960 and 1970 (as measured by income growth). Conversely, Kolko (2007) finds that poor neighborhoods bordering richer neighborhoods in 1990 grew faster between 1990 and 2000 than otherwise similar poor neighborhoods that were next to other poor neighborhoods. Our addition to this literature is that we propose a model that explains both of these facts and then formally test the model's predictions. During periods of declining city-wide housing demand in urban areas (like the suburbanization movement

during the 1960s), the richer neighborhoods on border of the rich areas will be the first to contract. Conversely, during periods of positive increases in city-wide housing demand (like that associated with the migration back to cities during the 1990s), the poor neighborhoods bordering the richer neighborhoods will be the first to gentrify.

Finally, recent work has focused on the effect of direct public policies on neighborhood gentrification. Specifically, Rossi-Hansberg et al. (2010) studies urban revitalization programs in Richmond, Virginia on neighborhood house price changes. Busso and Kline (2007) examine the effect of the federal urban Empowerment Zone program on neighborhood characteristics. Zheng and Kahn (2011) document the link between government investment and neighborhood gentrification within Beijing, China. Kahn et al. (2009) use differential zoning regulations along the California coast to examine gentrification patterns within and outside of the zoned areas. Our work complements this literature by highlighting gentrification that is not the result of government policy but instead endogenously results from the actions of private agents responding to city-wide housing demand shocks. Also, as we discuss below, this literature support our assumption that consumption externalities may be important in determining house prices within a city.

7.2.2 Importance of Consumption Externalities in Determining House Prices

A key mechanism in our model linking city-wide housing demand shocks, neighborhood gentrification, and neighborhood house price growth is the existence of neighborhood consumption externalities. Our theoretical model builds upon the insights of Benabou (1993) which looks at neighborhood sorting within a city where there are human capital externalities and the work of Becker and Murphy (2003) which looks at neighborhood sorting in a world with exogenous income groups where all agents have a preference to live around richer neighbors. From a theoretical standpoint, our work adds to this literature by examining the dynamics of sorting and house prices in response to city-wide housing demand shocks thereby generating a gentrification process.

Our work also adds to the growing literature showing, empirically, that cities are not only centers of production agglomeration, but also centers of consumption agglomeration. Glaeser et al. (2001) emphasize the importance of urban density in facilitating consumption. The important works of Bayer et al. (2004) and Bayer et al. (2007) empirically document neighborhood consumption externalities in that individuals are willing to pay more for educated and wealthier neighbors, all else equal. Autor et al. (2010) examine the importance of neighborhood spillovers during the removal of rent control within Cambridge, Massachusetts during last decade. As mentioned above, Rossi-Hansberg et al. (2010) estimate spillovers from urban revitalization programs implemented in Richmond, Virginia during the early 2000s.

Our work is also related to the sorting models of Tiebout (1956) where households of different income and preferences choose to locate in municipalities based on the provision of public goods provided by the municipalities. The municipally provided public goods are determined by the choices of residents in

the municipality.³² Our work is distinct from this literature in that we look at sorting within a given municipality where public expenditures and tax rates are determined at the municipality level.³³

7.2.3 Within City House Price Dynamics

There is very little systematic work documenting the variation of housing prices across neighborhoods within a city during housing price booms and busts. An important contribution of our work is to systematically document within-city patterns of house price appreciation during city-wide housing price booms. Aaronson (2001) documents a strong positive relationship between neighborhood house price growth and neighborhood income growth. Most of the existing papers that look at within-city house price dynamics compare the appreciation rate of “high end” properties to the appreciation rate of “low end” properties.³⁴ Two papers, however, do look specifically at differences in house price appreciation across zip codes within a metropolitan area. Case and Mayer (1996) look at differential movements in prices within cities of the Boston metro area between 1982 and 1992 while Case and Marynchenko (2002) look at differential trends in prices across different zip codes within the Boston, Chicago, and Los Angeles metro areas during the 1983 to 1993 period. No systematic relationships emerged from these studies. Our work extends this literature by systematically examining within-city movements in house prices during housing price booms. Moreover, we provide a framework that links such movements to city-wide housing demand shocks. In addition to our endogenous gentrification mechanism in response to city-wide housing demand shocks, our work also provides some facts that can be explored more fully in related papers that explain within-city house price dynamics. For a recent example building on our work, see Ferreira and Gyourko (2011).

7.3 Cross City House Price Dynamics

All of our empirical work focuses on how the endogenous gentrification mechanism can affect house price growth differentials across neighborhoods within a city. However, our model shows that the same mechanism can also generate differences in house price growth across cities either because of differences in the size of the shocks or because of differences in the underlying characteristics of the cities. On the one hand, cities that become richer endogenously gentrify, generating an increase in city-wide housing prices. On the other hand, richer cities may gentrify more in reaction to the same shock, experiences larger increase in city-wide housing prices.

Recently, Gyourko et al. (2006) highlight the importance of what they term “super star cities” to explain low frequency differences in house price growth across cities. These super-star cities are endowed

³²See recent work by Epple et al. (2001) and the cites within for a discussion of the theoretical extensions to the Tiebout model as well for corresponding empirical work.

³³There is also a large literature on residential sorting based on race. See, for example, Schelling (1969) and Card et al. (2008). All of our work controls for the racial composition of the neighborhood. Our work shows that income segregation is another important determinant of residential sorting above and beyond sorting based on race.

³⁴See, for example, Mayer (1993) and Poterba (1991) who analyzed Oakland, Dallas, Chicago, and Atlanta between 1970 and the mid 1980s, Case and Shiller (1994) who analyzed Boston and Los Angeles during the 1980s, and Smith and Tesarek (1991) who analyzed Houston during the 1970s and 1980s.

with a fixed desirable amenity (like good weather). Because proximity to the amenity is in fixed supply, households will bid up the land prices around the amenity as households become richer. In their model, there is a direct causation between household income within a city and house prices in certain cities that are endowed with desirable fixed amenities. This relationship, however, is pinned down because of the latent demand for a fixed amenity within the city. In our setup, the amenity is the rich people themselves. As higher income individuals migrate into the city, the nature of positive amenities provided by the city increases which puts upward pressure on house prices. While the two stories yield similar predictions, the mechanism driving the predictions are different.

Although we do not do so in this paper, we think it is an important area of future work to see how much of our endogenous gentrification mechanism can explain cross-city differences in house price appreciation rates.

8 Conclusion

In this paper, we explore the theoretical and empirical response of housing price dynamics across neighborhoods to a city-wide housing demand shock. As we show, positive city-wide housing demand shocks endogenously result in neighborhood gentrification. The key assumption in the model is that all individuals prefer neighborhoods populated by richer households as opposed to neighborhoods populated by poorer households. The reason for this is that richer neighborhoods endogenously provide amenities that are desirable to individuals. While we do not take a stand on the exact source of the externality, we have in mind that richer neighborhoods have lower levels of crime, higher provisions of local public goods, better peer effects, and a more extensive provision of service industries (like restaurants and entertainment options).

Empirically, we find that poor neighborhoods on the border of richer neighborhoods experience the largest increase in house price appreciation in response to a city-wide housing demand shock. In particular, we find that during the 1980s poor neighborhoods that bordered richer neighborhoods had house prices that appreciated by 6.8% more than otherwise similar poor neighborhoods which were farther away from rich neighborhoods in response to a one standard deviation Bartik shock. Moreover, these neighborhoods simultaneously experienced a more dramatic rise in resident income and education and a more dramatic decline in the resident poverty rate.

Moreover, our work adds to the work of Brueckner (1977) and Kolko (2007) by showing that there is a large spatial component to neighborhood gentrification. Aside from our results showing that it is the poor neighborhoods next to the rich neighborhoods that respond the most to city-wide housing demand shocks, we also show that proximity to rich neighborhoods is a defining feature of neighborhood gentrification. Choosing neighborhoods that have ex-post gentrified, we find that the probability of gentrification is 67 percent higher for neighborhoods that were within 0.5 miles of an existing rich neighborhood than

otherwise similar neighborhoods that were farther away. Our model and subsequent empirical work show that the spatial proximity of poor neighborhoods with respect to existing rich neighborhoods is a defining feature of neighborhoods that subsequently gentrify.

Aside from our dynamic model of neighborhood gentrification, we also bring new data to bear on the question. As far as we know, we are the first paper to systematically analyze the differential housing price dynamics across neighborhoods within cities during city-wide housing price booms. We show that the within-city variation in house prices is almost as large as the well documented cross-city variation in house prices during the last three decades. Also we document that poor neighborhoods experience larger housing price increases and a greater variation in housing price increases relative to richer neighborhoods during city-wide housing price booms. The larger the city-wide housing price boom, the more poor neighborhoods appreciate relative to rich neighborhoods. These patterns are robust to using different housing price indices and are found during all the time periods we analyze. Although our gentrification results only exploit the variation in house price appreciation among poor neighborhoods during city-wide housing price booms, the data we document suggests there are many other interesting patterns in the data worth exploring in future work.

The results in our paper can potentially explain the low frequency differences in housing price growth across cities during the last few decades. Part of the reason that San Francisco and Boston have experienced higher average annual house price growth during the last thirty years relative to cities like Des Moines and Nashville is that San Francisco and Boston may have gentrified more than than Des Moines and Nashville. Also, one area of future work would be to explore whether our mechanism can also explain short run movements in housing prices within and across cities. Moreover, it would be useful to incorporate uncertainty into the model. Can the expectation of a future positive housing demand shock coupled with an ex-post realization of a lower (negative) housing demand shock generate housing price booms and busts without any actual change in neighborhood characteristics? The expectation of future gentrification can cause prices to move today even if no actual gentrification takes place today. Exploring the theoretical and empirical predictions of such a model may yield additional important insights into the nature of housing price dynamics.

References

- Aaronson, D.**, “Neighborhood dynamics,” *Journal of Urban Economics*, 2001, 49 (1), 1–31.
- Alonso, William**, *Location and Land Use*, Harvard University Press, 1964.
- Autor, David H., Christopher J. Palmer, and Parag A. Pathak**, “Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge Massachusetts,” 2010.
- Bartik, T.J.**, “Who Benefits from State and Local Economic Development Policies?,” *W.E. Upjohn Institute for Employment Research: Kalamazoo, Mich.*, 1991.
- Bayer, P., R. McMillan, and K. Rueben**, “An equilibrium model of sorting in an urban housing market,” 2004.
- Bayer, Patrick, Robert McMillan, and Fernando Ferreira**, “A Unified Framework for Measuring Preferences for Schools and Neighborhoods,” *The Journal of Political Economy*, 2007, 115 (4), 588–638.
- Becker, Gary and Kevin Murphy**, *Social Economics: Market Behavior in a Social Environment*, Harvard University Press, 2003.
- Benabou, R.**, “Workings of a city: location, education, and production,” *The Quarterly Journal of Economics*, 1993, 108 (3), 619.
- Blanchard, O.J. and L.F. Katz**, “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1992, pp. 1–75.
- Brueckner, Jan K.**, “The Determinants of Residential Succession,” *Journal of Urban Economics*, 1977, 4 (1), 45–59.
- and **Stuart S. Rosenthal**, “Gentrification and Neighborhood Housing Cycles: Will America’s Future Downtowns be Rich?,” 2008.
- Busso, Matias and Patrick Kline**, “Do Local Economic Development Programs Work? Evidence from the Federal Empowerment Zone Program,” Unpublished 2007.
- Card, David, Alexandre Mas, and Jesse Rothstein**, “Tipping and The Dynamics of Segregation,” 2008, pp. 177–218.
- Case, Karl E. and Christopher J. Mayer**, “Housing Price Dynamics within a Metropolitan Area,” *Regional Science and Urban Economics*, 1996, 26 (3-4), 387–407.
- and **Maryna Marynchenko**, “Home Price Appreciation in Low-and Moderate-Income Markets,” *Low-income Homeownership: Examining the Unexamined Goal*, 2002, pp. 239–256.
- and **Robert J. Shiller**, “A Decade of Boom and Bust in the Prices of Single-Family Homes: Boston and Los Angeles, 1983 to 1993,” *New England Economic Review*, 1994, pp. 40–40.
- Davis, Morris A. and Francois Ortalo-Magné**, “Household Expenditures, Wages, Rents,” *Review of Economic Dynamics*, 2010.
- , — , and **Peter Rupert**, “What’s Really Going on in Housing Markets?,” Federal Reserve Bank of Cleveland, Economic Commentary 2007.
- Diamond, Rebecca**, “Estimation of City Amenity Valuations and Quality-of-Life with Preference Heterogeneity and Endogenous Sorting,” 2010.
- Epple, D., T. Romer, and H. Sieg**, “Interjurisdictional sorting and majority rule: An empirical analysis,” *Econometrica*, 2001, 69 (6), 1437–1465.
- Ferreira, Fernando and Joseph Gyourko**, “Anatomy of the Housing Boom in U.S. Neighborhoods and Metropolitan Areas, 1993-2009,” Unpublished Manuscript 2011.
- Glaeser, Edward L., Jed Kolko, and Albert Saiz**, “Consumer city,” *Journal of Economic Geography*, 2001, 1 (1), 27.
- , **Joseph Gyourko, and Charles Nathanson**, “Housing Dynamics,” Updated Version of NBER Working Paper 12787 2008.

- Gyourko, Joseph, Christopher J. Mayer, and Todd M. Sinai**, “Superstar Cities,” NBER Working Paper 12355 2006.
- Jr., Robert E. Lucas and Esteban Rossi-Hansberg**, “On the Internal Structure of Cities,” *Econometrica*, 2002, pp. 1445–1476.
- Kahn, Matthew E., Ryan Vaughn, and Jonathan Zasloff**, “The Housing Market Effects of Discrete Land Use Regulations: Evidence from the California Coastal Boundary Zone,” 2009.
- Kolko, Jed**, “The Determinants of Gentrification,” Unpublished manuscript 2007.
- , “Job Location, Neighborhood Change, and Gentrification,” Unpublished manuscript 2009.
- Mayer, Christopher J.**, “Taxes, Income Distribution, and the Real Estate Cycle: Why all Houses do not Appreciate at the Same Rate,” *New England Economic Review*, 1993, pp. 39–39.
- McKinnish, T., R. Walsh, and T. Kirk White**, “Who Gentrifies Low-Income Neighborhoods?,” *Journal of urban economics*, 2010, 67 (2), 180–193.
- Mian, Atif and Amir Sufi**, “The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis,” *Quarterly Journal of Economics*, 2009, 124 (4).
- Mills, Edwin S.**, “An Aggregative Model of Resource Allocation in a Metropolitan Area,” *The American Economic Review*, 1967, pp. 197–210.
- Muth, Richard F.**, *Cities and Housing*, Univ. of Chicago Press, 1969.
- Nieuwerburgh, S. Van and P.O. Weill**, “Why Has House Price Dispersion Gone Up?,” *The Review of Economic Studies*, 2010.
- Notowidigdo, M.J.**, “The Incidence of Local Labor Demand Shocks,” 2010.
- Poterba, James M.**, “House Price Dynamics: The Role of Tax Policy and Demography,” *Brookings Papers on Economic Activity*, 1991, pp. 143–203.
- Roback, Jennifer**, “Wages, Rents, and the Quality of Life,” *The Journal of Political Economy*, 1982, 90 (6), 1257.
- Rosen, Sherwin**, “Wage-based Indexes of Urban Quality of Life,” *Current Issues in Urban Economics*, 1979, 3.
- Rossi-Hansberg, Esteban, Pierre-Daniel Sartre, and Raymond Owens III**, “Housing Externalities,” *Journal of Political Economy*, 2010, 118 (3), 485–535.
- Saiz, A.**, “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics*, 2010.
- Schelling, T.C.**, “Models of Segregation,” *The American Economic Review*, 1969, 59 (2), 488–493.
- Smith, Barton A. and William P. Tesarek**, “House Prices and Regional Real Estate Cycles: Market Adjustments in Houston,” *Real Estate Economics*, 1991, 19 (3), 396–416.
- Tiebout, C.M.**, “A pure theory of local expenditures,” *The journal of political economy*, 1956, pp. 416–424.
- Vigdor, J.L., D.S. Massey, and A.M. Rivlin**, “Does Gentrification Harm the Poor?[with Comments],” *Brookings-Wharton Papers on Urban Affairs*, 2002, pp. 133–182.
- Zheng, S. and M.E. Kahn**, “Does Government Investment in Local Public Goods Spur Gentrification? Evidence from Beijing,” NBER Working Paper 2011.

Housing Price Growth versus Initial Housing Price 2000-2006: New York MSA




Figure 1: Figure shows the initial house price in a zip code in 2000 versus the subsequent house price growth in that zip code between 2000 and 2006. The sample includes all zip codes within the New York metro area for which a Case-Shiller house price index exists. We measure the initial house price using median home value from the 2000 Census. We use the Case-Shiller index to compute the growth rate in house price between 2000 and 2006.

Model Generated Externality and House Prices across Neighborhoods




Figure 2: Figure shows the model predicted relationship between the size of the rich neighborhood (top panel), the value of the neighborhood externality (middle panel), and the house price in the neighborhood (bottom panel).




Figure 3: We set $\alpha = .8, \beta = .8, \delta^R = .2, \delta^P = 0, A = 1, \gamma = .1, r = .03, y^R = 1, y^P = .5, C^R = C^P = 25, N^R = N^P = .5$. The shock is an unexpected and permanent increase in ϕ from $\phi = 1$ to $\phi = 5$.

Table 1: House Price Growth Variation Across MSAs and Across Neighborhoods Within an MSA or City

Time Period	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
2000 - 2006 (observations)	0.33 (384)	0.42 (20)	0.18 (1,617)	0.18 (471)	0.24 (471)	-	-	-
1990 - 2000 (observations)	0.17 (348)	0.21 (17)	0.15 (1,512)	0.17 (496)	-	0.15 (496)	0.33 (9,684)	0.54 (16,161)
1980 - 1990 (observations)	0.31 (158)	-	-	-	-	-	0.24 (4,640)	0.44 (8,729)
Variation	Between MSA	Between MSA	Within MSA	Within City	Within City	Within City	Within City	Within City
Neighborhood Definition	-	-	Zip Code	Zip Code	Zip Code	Zip Code	Census Tract	Census Tract
House Price Data	FHFA	Case Shiller	Case Shiller	Case Shiller	Zillow	Census	Census	Census

Notes: Table shows the between MSA standard deviation of house prices (columns 1 and 2) and the within-city/MSA standard deviation across neighborhoods (remaining columns) for different house price measures, different time periods, and different definitions of neighborhoods. The Case-Shiller data is only available for the 1990s and the 2000s. The Zillow data is available only in the 2000s. The Census house price growth measure is available only during the 1980s and 1990s. For column 1, we use all 384 MSAs available in the FHFA data. For column 2, we use the 20 MSAs for which Case-Shiller reports an index. For columns 3 - 6, we use all available zip codes for which a reliable Case-Shiller index exists. See the text for details. In column 7, we use all census tracts in Case-Shiller main cities. In column 8, we use all census tracts in all cities where there at least 30 census tracts within the city. See text for additional details. For the Census data, the top and bottom 1 percent of neighborhoods with respect to median home price growth are dropped.

Table 2: Housing Price Growth by Initial Price Quartile

	(1) Quartile 4	(2) Quartile 3	(3) Quartile 2	(4) Quartile 1	(5) p-val of Quartile 4 = Quartile 1
<u>2000 - 2006: Housing Booms</u>					
Chicago, City Level, Case-Shiller	0.53	0.66	0.72	0.88	0.00
Chicago, MSA Level, Case-Shiller	0.47	0.50	0.49	0.69	0.00
New York City, MSA Level, Case-Shiller	0.64	0.76	0.86	1.11	0.00
Boston, MSA Level, Case-Shiller	0.40	0.47	0.54	0.62	0.00
Los Angeles, MSA Level, Case-Shiller	1.21	1.40	1.58	1.76	0.00
San Francisco, MSA Level, Case-Shiller	0.35	0.41	0.49	0.61	0.00
Washington D.C., MSA Level, Case-Shiller	1.29	1.37	1.49	1.61	0.00
<u>1990 - 1997: Housing Booms</u>					
Denver, MSA Level, Case-Shiller	0.51	0.50	0.52	0.89	0.00
Portland, MSA Level, Case-Shiller	0.41	0.52	0.49	0.69	0.00
<u>1984 - 1989: Housing Booms</u>					
New York City, City Level, Furman	0.33	0.57	0.69	1.06	0.00
Boston, MSA Level, Case-Shiller	0.65	0.69	0.75	0.84	0.00
<u>1990 - 1997: Housing Busts</u>					
San Francisco, MSA Level, Case-Shiller	-0.08	-0.07	-0.11	-0.14	0.00
Boston, MSA Level, Case-Shiller	-0.05	-0.08	-0.13	-0.11	0.00

Notes: This table shows the mean house price appreciation rates for neighborhoods grouped by quartile of initial housing prices, during different time periods, and using different housing price indices to measure the price appreciation. Quartile 4 has the highest initial price zip codes within the city while quartile 1 has the lowest initial price zip codes within the city. Each row labels a city or metro area for a given time period using a given house price index to compute the appreciation rates.

Table 3: Regression of Neighborhood House Price Appreciation on Distance to Nearest High-Price Neighborhood and Other Controls, Across Different Samples With Different House Price Measures

Time Period	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Distance to Nearest High-Price Neighborhood	-0.074 (0.021)	-0.064 (0.020)	-0.059 (0.032)	-0.043 (0.030)	-0.071 (0.036)	-0.234 (0.042)	-0.138 (0.031)	-0.139 (0.028)
Log Distance to Nearest High-Price Neighborhood * City Wide Bust Indicator	-	-	-	-	0.084 (0.048)	0.070 (0.054)	0.078 (0.029)	0.083 (0.031)
House Price Measure/ Neighborhood Aggregation	C-S Zip Code	C-S Zip Code	Zillow Zip Code	C-S Zip Code	C-S Zip Code	Census Census Tract	Census Census Tract	Census Census Tract
Time Period	00-06	00-06	00-06	90-00	90-00	90-00	90-00	80-90
Vector of Z Controls Included	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of Observations	211	211	211	223	223	3,015	7,981	4,251
Mean Log Distance to Nearest High-Price Neighborhood	1.23	1.23	1.23	1.22	1.22	0.397	0.497	0.320
Std. Dev. Log Distance to Nearest High-Price Neighborhood	0.538	0.538	0.538	0.488	0.488	0.784	0.719	0.717

Note: Table shows regression of neighborhood house price appreciation between period t and $t+k$ on log distance to nearest high price neighborhood within the neighborhood's city, city fixed effects, and a vector of neighborhood controls. High price neighborhoods are those neighborhoods that are within the top quartile of average neighborhood house prices in year t . We restrict our analysis in this table to those neighborhoods within the city which were in the bottom half of the house price distribution in period t . See text for additional sample descriptions and discussion of the controls included. Robust standard errors, clustered by MSA, are shown in parentheses. All regressions are weighted by the number of owner occupied housing units in the neighborhood in the initial year.

Table 4: Instrumental Variable Regression of House Price Appreciation on Proximity to Nearest High-Price Neighborhood, Census Data 1980 - 1990

	(1)	(2)
Log Dist. to Nearest High-Price Neighborhood	0.34 (0.10)	-
Log Dist. to Nearest High-Price Neighborhood * Bartik Predicted City-Wide Income Shock	-2.27 (0.53)	-
Dummy: High-Price Neighborhood Within 0 - 1 Miles	-	-0.63 (0.24)
Dummy: High-Price Neighborhood Within 1 - 3 Miles	-	-0.15 (0.11)
Dummy: High-Price Neighborhood Within 0 - 1 Miles * Bartik Predicted City-Wide Income Shock	-	4.06 (1.27)
Dummy: High-Price Neighborhood Within 1 - 3 Miles * Bartik Predicted City-Wide Income Shock	-	1.00 (0.58)

Note: Table reports the results from the regression specified by Equation (22) from the text. The sample is the same as that used in Table 3, column 8. Standard errors bootstrapped. First stage is re-sampled from IPUMS, second stage from census tract tabulations, and stratified by city (2,500 repetitions). P-value on test of whether last two coefficients in column 2 are equal is 0.020.

Table 5: Regression of Neighborhood House Price Appreciation on Neighborhood Income Growth and Other Controls, Across Different Samples With Different House Price Measures

	(1)	(2)	(3)	(4)	(5)
Neighborhood Income Growth	0.677 (0.306)	0.625 (0.176)	0.452 (0.096)	0.326 (0.052)	0.197 (0.043)
House Price Measure	C-S	Census	Census	Census	Census
Neighborhood Aggregation	Zip Code	Zip Code	Census Tract	Census Tract	Census Tract
Time Period	90-00	90-00	90-00	90-00	80-90
Number of Observations	223	223	3,015	7,981	4,251
Mean Neighborhood Income Growth	0.069	0.069	0.080	0.100	0.113
Std. Dev. Neighborhood Income Growth	0.151	0.151	0.273	0.258	0.284

Note: Table shows regression of neighborhood house price appreciation between period t and $t+k$ on neighborhood income growth, city fixed effects, and a vector of neighborhood controls. As in Table 3 and 4, we restrict our analysis to those neighborhoods in the bottom half of the neighborhood house price distribution in period t . The specifications in columns 1 and 2 use low price zip codes from Case-Shiller cities where a Case-Shiller price index exists in 1990 and 2000. The specification in column 3 uses all census tracts in Case-Shiller cities. The specifications in column 4 uses all census tracts from all cities in the U.S. which have at least 30 consistently defined census tracts between 1990 and 2000. The specification in column 5 uses all census tracts from all cities in the U.S. which have at least 30 consistently defined census tracts between 1980 and 1990. For the specifications in columns 3 - 5, we also trim the top and bottom 1 percent of the house price growth and the income growth distributions. See text for additional details. Robust standard errors clustered at the city level are in parentheses.

Table 6: Measures of Gentrification and Proximity to High Income Neighborhoods: Descriptive Analysis 1980 - 1990 and 1990 - 2000

	Gentrification Measure:			
	Neighborhood Income Growth During Time Period			
	Greater than 50%		Greater than 25%	
	(1)	(2)	(3)	(4)
	1980 - 1990	1990 - 2000	1980 - 1990	1990 - 2000
Dummy: High-Price Neighborhood Within 0 - 0.5 Miles	0.069 (0.017)	0.057 (0.027)	0.082 (0.035)	0.109 (0.040)
Dummy: High-Price Neighborhood Within 0.5 - 1 Miles	0.015 (0.007)	0.017 (0.009)	0.092 (0.020)	0.062 (0.020)
Dummy: High-Price Neighborhood Within 1 - 2 Miles	0.006 (0.008)	0.018 (0.007)	0.076 (0.020)	0.029 (0.014)
Dummy: High-Price Neighborhood Within 2 - 3 Miles	-0.005 (0.007)	0.002 (0.005)	0.024 (0.019)	0.018 (0.014)
Fraction of Neighborhoods that Gentrified	11.0%	5.9%	30.2%	19.7%
Sample Size	4,251	7,981	4,251	7,981

Note: Table shows the results from a linear probability regression of a dummy variable indication whether a neighborhood gentrified between t and $t+k$ on the proximity of that neighborhood to an existing rich neighborhood. We define rich neighborhoods as those census tracts within a city that are in the top quartile of the period t house price distribution. The samples in columns 1 and 3 (columns 2 and 4) are the same as column 8 (column 7) of Table 3. All regressions include city fixed effects. Robust standard errors, clustered at the city level, are in parentheses.

Table 7: Instrumental Variable Estimation of Measures of Gentrification and Proximity to High Income Neighborhoods, Census Data 1980 - 1990

	Dependent Variable: Measure of Neighborhood Gentrification		
	(1) Growth in Median Income	(2) Change in Poverty Rate	(3) Change in Fraction of Residents with Bachelor's Degree
Log Dist. to Nearest High-Price Neighborhood	0.10 (0.05)	-0.04 (0.02)	0.04 (0.01)
Log Dist. to Nearest High-Price Neighborhood *	-0.57 (0.27)	0.22 (0.12)	-0.23 (0.08)
Bartik Predicted City-Wide Income Shock			
Mean: Dependent Variable	0.149	0.029	0.028
Sample Size	4,251	4,251	4,251

Note: Table reports the results from the regression specified by Equation (23) from the text. The sample is the same as that used in Table 3, column 8. Standard errors bootstrapped. First stage is re-sampled from IPUMS, second stage from census tract tabulations, and stratified by city (2,500 repetitions).

Table A1: Correlation Between Housing Price Indices

Sample	Correlation Between Case-Shiller and Zillow Appreciation 2000 - 2006	Correlation Between Case-Shiller and Census Appreciation 1990 - 2000
All Case-Shiller Zip Codes	0.96 (n = 1,617)	0.78 (n = 1,488)
Case-Shiller Zip Codes from Main Cities	0.96 (n = 471)	0.82 (n = 496)

Notes: Column 1 of this table shows the correlation across zip codes between the percentage change in the Case-Shiller house price index and the percentage change in the Zillow house price index between 2000 and 2006. Row 1 restricts the sample to all zip codes for which both a Zillow and Case-Shiller zip code price index is available. Row 2 restricts the sample only to those available zip codes within the main city of the Case-Shiller MSAs. See text for details. The second column shows the correlation across zip codes between the percentage change in the Case-Shiller zip code level housing price index and the percentage change in the median reported house price within the zip code from the U.S. Census between 1990 and 2000. Sample sizes differ slightly between the two columns because the Census median house price and the Zillow price index are not available for all Case-Shiller zip codes.

Table A2: Regression of Percentage Change in Case-Shiller Index on Percentage Change in Zillow and Census Price Indices

	Independent Variable			
	Zillow Price Appreciation 2000 - 2006		Census Median Price Appreciation 1990 - 2000	
	(1)	(2)	(3)	(4)
Coefficient	1.06 (0.01)	1.02 (0.02)	0.96 (0.03)	1.02 (0.06)
Constant	0.04 (0.01)	0.09 (0.01)	0.02 (0.02)	0.07 (0.03)
R-squared	0.92	0.92	0.66	0.71
Sample	All	Main City	All	Main City
Sample Size	1,617	471	1,487	496

Notes: Columns 1 and 2 of this table show the simple regression of the percentage change in the Case-Shiller index for a given neighborhood within a city between 2000 and 2006 on the percentage change in the Zillow index for the same neighborhood between 2000 and 2006. In columns 3 and 4, the table shows the simple regression of the percentage change in the Case-Shiller index between 1990 and 2000 on the percentage change in median house prices from the Census between 1990 and 2000 and controls for the change in neighborhood structure characteristics from the Census during the same time period. The change in neighborhoods structure characteristics include: the change in the fraction of housing units that are single family detached houses, the change in the fraction have 0 or 1 bedrooms, the change in the fraction that have 2 bedrooms, the change in the fraction that have 3 bedrooms, the change in the fraction of housing units that were built in the past 5 years, the change in the fraction that were built 5 - 20 years ago, the change in the fraction that were built 20 - 40 years ago, and the change in the fraction that were built 40 - 50 years ago. Our measure of neighborhood is zip code. The sample includes all zip codes where both measures of price increases exist. All regressions are weighted by the number of owner occupied housing units in the zip code. The differences between columns 1 and 2 (or analogously 3 and 4) are that column 1 includes all zip codes within the Case-Shiller MSAs while column 2 only includes the zip codes from the main city within the Case-Shiller MSAs.

Table A3: MSA Level House Price Appreciation for Case-Shiller MSAs

	In MSA Sample, City Sample, or Both	MSA Price Appreciation 2000 - 2006	MSA Price Appreciation 1990 - 2000
Akron	Both	3.6%	22.8%
Atlanta	Both	13.8	14.6
Boston	MSA Only	49.0	12.5
Charlotte	Both	9.2	12.7
Chicago	Both	36.8	14.4
Cincinnati	Both	7.4	15.4
Columbus (OH)	Both	7.4	16.8
Denver	Both	10.6	65.9
Fresno	Both	124.1	-8.5
Jacksonville	Both	69.4	11.2
Las Vegas	Both	88.8	-0.3
Los Angeles	Both	121.7	-20.9
Memphis	Both	6.0	8.7
Miami	Both	125.6	15.3
New York	Both	72.4	0.0
Oakland	Both	76.7	9.0
Philadelphia	Both	59.1	-10.4
Phoenix	Both	82.8	20.8
Portland (OR)	Both	47.0	53.4
Raleigh	Both	8.1	14.5
Sacramento	Both	96.0	-11.7
San Francisco	MSA Only	52.1	18.9
San Diego	Both	93.5	-0.7
San Jose	Both	44.0	30.6
Seattle	Both	46.7	22.9
St. Paul	Both	38.2	29.9
Tampa	Both	88.6	7.0
Toledo	Both	4.7	19.7
Washington DC	MSA Only	98.3	-8.5

Notes: Table shows the MSA level house price appreciation rates for each MSA for the 2000-2006 period (column 2) and the 1990-2000 period (column 3) using the FHFA MSA level house price indices. Column 1 is an indicator whether the data from these cities or MSAs are included in our MSA samples or in our main city only samples. The data from Boston, San Francisco, and Washington DC are not included in our main city sample because there are not 10 zip codes within the city that has a reliably computed price index.

A1 Theory Appendix (Included in Publication)

A1.1 Proof of Proposition 1

Most of the proof of Proposition 1 is in the text. As we argue in the text, we are left to check only that

$$\begin{aligned} U^R(i) &\leq \bar{U}^R \text{ for all } i \in [I_t, \bar{I}_t], \\ U^P(i) &\leq \bar{U}^P \text{ for all } i \in [0, I_t], \end{aligned}$$

where $U^s(i)$ is defined in expression (6). Using expression (7), these two conditions can be rewritten as

$$K^R(A + H_t(i))^{\frac{\delta_R}{\beta}} \leq K^P(A + H_t(i))^{\frac{\delta_P}{\beta}} + \frac{r}{1+r}(C^R - C^P) \text{ for all } i \in [I_t, \bar{I}_t], \quad (22)$$

$$K^P(A + H_t(i))^{\frac{\delta_P}{\beta}} \leq K^R(A + H_t(i))^{\frac{\delta_R}{\beta}} - \frac{r}{1+r}(C^R - C^P) \text{ for all } i \in [0, I_t]. \quad (23)$$

Combining (9) with (13) and (14) we obtain

$$\begin{aligned} K^R &= \frac{r}{1+r} \left[C^P \left(\frac{A}{A+\gamma} \right)^{-\frac{\delta_P}{\beta}} + (C^R - C^P) \right] (A + \gamma)^{-\frac{\delta_R}{\beta}} \\ K^P &= \frac{r}{1+r} C^P A^{-\frac{\delta_P}{\beta}}, \end{aligned}$$

Using these expressions, condition (22) can be rewritten as

$$\left(\frac{A + H_t(i)}{A + \gamma} \right)^{\frac{\delta_R - \delta_P}{\beta}} \leq \frac{1 + \left(\frac{C^R - C^P}{C^P} \right) \left(\frac{A}{A + H_t(i)} \right)^{\frac{\delta_P}{\beta}}}{1 + \left(\frac{C^R - C^P}{C^P} \right) \left(\frac{A}{A + \gamma} \right)^{\frac{\delta_P}{\beta}}}.$$

for all $i \in [I_t, \bar{I}_t]$. This implies that $H_t(i) < \gamma$ and hence the RHS is not smaller than 1 and that, if $\delta_R \geq \delta_P$, the LHS is not bigger than 1. Hence, $\delta_R \geq \delta_P$ is a sufficient condition for this condition to be satisfied. Notice that if $C^R = C^P$, this is also a necessary condition.

Next, condition (23) can be rewritten as

$$\left(\frac{A + H_t(i)}{A + \gamma} \right)^{\frac{\delta_P - \delta_R}{\beta}} \leq 1 + \left(\frac{C^R - C^P}{C^P} \right) \left(\frac{A}{A + \gamma} \right)^{\frac{\delta_P}{\beta}} \left[1 - \left(\frac{A + \gamma}{A + H_t(i)} \right)^{\frac{\delta_R}{\beta}} \right]$$

for all $i \in [0, I_t]$. In these locations, by construction, $H_t(i) > \gamma$, which implies that the RHS is not smaller than 1 and that, if $\delta_R \geq \delta_P$ the LHS is not bigger than 1. It follows that $\delta_R \geq \delta_P$ is also a sufficient condition for this equation to hold. Again, it is also a necessary condition if $C^R = C^P$. Hence, this completes the proof that a fully segregated equilibrium exists if $\delta^P \leq \delta^R$.

A1.2 Proof of Proposition 2

The initial price schedule is:

$$p_t(i) = \begin{cases} p_t^R(i) & \text{for } i \in [0, I_t] \\ p_t^P(i) & \text{for } i \in [I_t, \bar{I}_t] \end{cases}, \quad (24)$$

where $p_t^R(i)$ and $p_t^P(i)$ are given by (17) and (18). First, notice that if $i \geq I_t + \gamma$, then $p_t(i) = C^P$, and if $i < I_t + \gamma$, then $p_t(i) > C^P$. Also, if $i < I_t - \gamma$, then $p_t(i) = \bar{p}$, where

$$\bar{p} \equiv \left[C^P \left(1 + \frac{\gamma}{A} \right)^{\frac{\delta_P}{\beta}} + C^R - C^P \right] \left(1 + \frac{\gamma}{A + \gamma} \right)^{\frac{\delta_R}{\beta}}.$$

Now, imagine that the economy is hit by an unexpected and permanent increase in population, so that the measure of both rich and poor households increase by a proportion of $\phi > 1$, *i.e.* $N_{t+1}^s = \phi N_t^s$ for $s = R, P$. We then have

$$\frac{p_{t+1}(i)}{p_t(i)} = \begin{cases} \left(\frac{A+\gamma+\min\{\gamma, I_{t+1}-i\}}{A+\gamma+\min\{\gamma, I_t-i\}} \right)^{\frac{\delta_R}{\beta}} & \text{for } i \in [0, I_t] \\ \left[\left(\frac{A+\gamma}{A} \right)^{\frac{\delta_P}{\beta}} + \frac{C^R - C^P}{C^P} \right] \frac{\left(1 + \frac{\min\{\gamma, I_{t+1}-i\}}{A+\gamma} \right)^{\frac{\delta_R}{\beta}}}{\left(1 + \frac{\max\{\gamma+I_t-i, 0\}}{A} \right)^{\frac{\delta_P}{\beta}}} & \text{for } i \in [I_t, I_{t+1}] \\ \left(\frac{A+\max\{\gamma+I_{t+1}-i, 0\}}{A+\max\{\gamma+I_t-i, 0\}} \right)^{\frac{\delta_P}{\beta}} & \text{for } i \in [I_{t+1}, \bar{I}_t] \end{cases} \quad (25)$$

Also, from equations (15) and (16), we obtain $I_{t+1} > I_t$ and $\bar{I}_{t+1} > \bar{I}_t$. Then, if $i < I_t - \gamma$, it must be that $p_{t+1}(i)/p_t(i) = 1$, which implies that

$$E_{t+1} \left[\frac{p_{t+1}(i)}{p_t(i)} \mid p_t(i) = \bar{p} \right] = 1.$$

Moreover, $I_{t+1} > I_t$, together with expression (25), immediately implies that $p_{t+1}(i)/p_t(i) \geq 1$ for $i > I_t - \gamma$, and hence

$$E_{t+1} \left[\frac{p_{t+1}(i)}{p_t(i)} \mid p_t(i) < \bar{p} \right] > 1,$$

which proves the first statement of the proposition.

We now want to prove the second statement of the proposition, that is, that the price ratio $p_{t+1}(i)/p_t(i)$ is non-increasing in $p_t(i)$. First, notice that $p_t(i)$ is non-increasing in i , so proving that $p_{t+1}(i)/p_t(i)$ is non-increasing in $p_t(i)$ is equivalent to prove that $p_{t+1}(i)/p_t(i)$ is non-decreasing in i . The ratio $p_{t+1}(i)/p_t(i)$ is continuous and differentiable except at a finite number of points. Hence, in order to prove that it is non-decreasing in i , it is enough to show that $d[p_{t+1}(i)/p_t(i)]/di$ is non-negative, for all i where this derivative exists. Let us show that.

For $i \in [0, I_t - \gamma]$, $p_{t+1}(i)/p_t(i) = 1$ and hence $p_{t+1}(i)/p_t(i)$ is constant in i . For $i \in [I_t - \gamma, I_t]$, we have that

1. if $I_t - \gamma < i < I_{t+1} - \gamma$, then

$$\frac{d\left(\frac{p_{t+1}(i)}{p_t(i)}\right)}{di} = \frac{\delta_R}{\beta(A+\gamma)} \left(\frac{A+2\gamma}{A+\gamma} \right)^{\frac{\delta_R}{\beta}} \left(\frac{A+\gamma+I_t-i}{A+\gamma} \right)^{-\frac{\delta_R}{\beta}-1} > 0 \quad (26)$$

2. if $I_t - \gamma < i < \min\{I_{t+1} - \gamma, I_t\}$, then

$$\frac{d\left(\frac{p_{t+1}(i)}{p_t(i)}\right)}{di} = \frac{\delta_R}{\beta} \left(\frac{A+\gamma+I_{t+1}-i}{A+\gamma+I_t-i} \right)^{\frac{\delta_R}{\beta}} \left[\frac{1}{A+\gamma+I_t-i} - \frac{1}{A+\gamma+I_{t+1}-i} \right] > 0 \quad (27)$$

given that $I_t < I_{t+1}$.

For $i \in [I_t, I_t + \gamma]$ we have that

1. if $I_t < i < \min\{I_t + \gamma, I_{t+1} - \gamma\}$

$$\frac{d\left(\frac{p_{t+1}(i)}{p_t(i)}\right)}{di} = \frac{\tilde{C}}{A} \left(\frac{A+2\gamma}{A+\gamma} \right)^{\frac{\delta_R}{\beta}} \left(\frac{A+\gamma+I_t-i}{A} \right)^{-\frac{\delta_P}{\beta}-1} > 0 \quad (28)$$

where

$$\tilde{C} \equiv \left[\left(\frac{A+\gamma}{A} \right)^{\frac{\delta_P}{\beta}} + \frac{C^R - C^P}{C^P} \right]$$

2. if $\max \{I_t, I_{t+1} - \gamma\} < i < \min \{I_{t+1}, I_t - \gamma\}$

$$\frac{d \left(\frac{p_{t+1}(i)}{p_t(i)} \right)}{di} = \frac{\tilde{C} \left(\frac{A+\gamma+I_{t+1}-i}{A+\gamma} \right)^{\frac{\delta_R}{\beta}}}{\left(\frac{A+\gamma+I_t-i}{A} \right)^{\frac{\delta_P}{\beta}}} \left[\frac{\delta_P}{A+\gamma+I_t-i} - \frac{\delta_R}{A+\gamma+I_{t+1}-i} \right] \quad (29)$$

hence

$$\frac{d \left(\frac{p_{t+1}(i)}{p_t(i)} \right)}{di} > 0 \text{ iff } \frac{\delta_R}{\delta_P} < \frac{A+\gamma+I_{t+1}-i}{A+\gamma+I_t-i},$$

which is true if the shock is big enough and $I_{t+1} - I_t$ is big enough;

3. if $\max \{I_t, I_{t+1}\} < i < I_t + \gamma$

$$\frac{d \left(\frac{p_{t+1}(i)}{p_t(i)} \right)}{di} = \frac{\delta_P}{\beta} \left(\frac{A+\gamma+I_{t+1}-i}{A+\gamma+I_t-i} \right)^{\frac{\delta_P}{\alpha}} \left[\frac{1}{A+\gamma+I_t-i} - \frac{1}{A+\gamma+I_{t+1}-i} \right] > 0. \quad (30)$$

This proves that, if the shock is big enough, the second statement of the proposition holds.

Proof of Proposition 3

The proof of this Proposition is straightforward. Imagine that at time $t+1$ the economy is hit by an unexpected and permanent increase in population, that is, $N_{t+1}^s = \phi N_t^s$ with $\phi > 1$ for $s = P, R$. From expressions (15) and (16) it follows that $I_{t+1} > I_t$ and $\bar{I}_{t+1} > \bar{I}_t$. Then, from expression (25), we immediately obtain that for all $i \leq I_t + \gamma$, that is, for all i such that $p_t(i) = C^P$, $d(p_{t+1}(i)/p_t(i))/di < 0$, as we wanted to show.

A1.3 Proof of Proposition 4

First, notice that at time t , each location i may lie in four possible intervals that implies different pricing behavior: $[0, I_t - \gamma]$, $[I_t - \gamma, I_t]$, $[I_t, I_t + \gamma]$, and $[I_t, \bar{I}_t]$. From expression (24), it is immediate that prices at time $t+1$ in each location i are weakly increasing in I_{t+1} , whenever i is in the same type of interval at t and $t+1$. From expression (15) where N^s is substituted by ϕN^s for $s = R, P$, I_{t+1} is non-decreasing in ϕ and hence prices are weakly increasing in ϕ for all i which remain in the same type of interval. Let us consider any $\phi^A > \phi^B > 1$, with $I_{t+1}^A > I_{t+1}^B$. Then all $i \in [0, I_{t+1}^A - \gamma]$ are also in $[0, I_{t+1}^B - \gamma]$, but some $i \in [I_{t+1}^B - \gamma, I_{t+1}^B + \gamma]$ may be in $[0, I_{t+1}^A - \gamma]$ or some $i \in [I_{t+1}^B, I_{t+1}^B + \gamma]$ may be in $[I_{t+1}^A - \gamma, I_{t+1}^A]$. Given that, from inspection of expression (24), $p_{t+1}(i)$ is non-increasing in i , this implies that aggregate prices P_{t+1} must be non-decreasing in ϕ . Hence, if at time $t+1$ the economy is hit by an unexpected and permanent increase in ϕ , then P_{t+1} is going to be higher, the larger is the increase in ϕ . Given that P_t is given, this immediately proves the first statement of the proposition that the percentage increase in aggregate price is higher the larger is the increase in ϕ .

Second, we want to prove the second statement of the proposition, that

$$\frac{d^2(p_{t+1}(i)/p_t(i))}{dp_t(i)d\phi} \geq 0$$

for all $p_t(i) > C^P$ where the derivative is well-defined. Equations (26)-(30) in the proof of Proposition (2) define $d[p_{t+1}(i)/p_t(i)]/di$ for all i where this derivative is well-defined and $p_t(i) > C^P$. If the increase in ϕ is big enough, $d[p_{t+1}(i)/p_t(i)]/di > 0$ for all $p_t(i) > C^P$. Moreover, by inspection, it is easy to see that $d[p_{t+1}(i)/p_t(i)]/di$ is increasing in I_{t+1} , and hence increasing in ϕ , whenever i is in the same type of interval after a small or a large shock, say ϕ^A or ϕ^B . Moreover, given that $I_{t+1}^A > I_{t+1}^B$, i may lie in different types of interval in the two cases. In particular, it could be that $\min \{I_{t+1}^B - \gamma, I_t^A\} < i < I_t^A$ but $I_t^B - \gamma < i < \min \{I_t^A - \gamma, I_t^B\}$, or that $\max \{I^H, I_B^L - \gamma\} < i < I^H + \gamma$ and $I^H < i < \min \{I^H + \gamma, I_A^L - \gamma\}$, or that $I_B^L < i < I^H + \gamma$ but $\max \{I^H, I_A^L - \gamma\} < i < I^H + \gamma$. It is easy to see that expression (26) is not smaller

than expression (27) and that expression (28) is not smaller than expression (29). Finally expression (29) is bigger than expression (30) iff

$$\left(\frac{A + \gamma + I_{t+1} - i}{A + \gamma} \right)^{\frac{\delta_R - \delta_P}{\beta}} \left[1 - \frac{(\delta_R - \delta_P)(A + \gamma + I_t - i)}{\delta_P(I_{t+1} - I_t)} \right] > 1,$$

which is true if the shock is large enough so that $I_{t+1} - I_t$ is big enough, as we assumed. This proves that $d^2 [p_{t+1}(i)/p_t(i)]/did\phi$ is positive for all i such that the derivative exists and $p_t(i) > C^P$. Given that $p_t(i)$ is non-increasing in i , this completes the proof of the second claim of the proposition.