NBER WORKING PAPER SERIES

DO PEOPLE SHAPE CITIES, OR DO CITIES SHAPE PEOPLE? THE CO-EVOLUTION OF PHYSICAL, SOCIAL, AND ECONOMIC CHANGE IN FIVE MAJOR U.S. CITIES

Nikhil Naik
Scott Duke Kominers
Ramesh Raskar
Edward L. Glaeser
César A. Hidalgo

Working Paper 21620
http://www.nber.org/papers/w21620

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2015

We would like to acknowledge helpful comments from Gary Becker, Jörn Boehnke, Steven Durlauf, James Evans, Jay Garlapati, Lars Hansen, James Heckman, John Eric Humphries, Jackie Hwang, Priya Ramaswamy, Robert Sampson, Zak Stone, Erik Strand, and Nina Tobio. Mia Petkova contributed Figure 2. N.N. acknowledges the support from The MIT Media Lab consortia; S.D.K. acknowledges support from the National Science Foundation (grants CCF-1216095 and SES-1459912), the Harvard Milton Fund, the Wu Fund for Big Data Analysis, and the Human Capital and Economic Opportunity Working Group (HCEO) sponsored by the Institute for New Economic Thinking (INET); E.L.G. acknowledges support from the Taubman Center for State and Local Government; and C.A.H. acknowledges support from Google's Living Lab awards and The MIT Media Lab consortia. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w21620.ack

Do People Shape Cities, or Do Cities Shape People? The Co-evolution of Physical, Social, and Economic Change in Five Major U.S. Cities
Nikhil Naik, Scott Duke Kominers, Ramesh Raskar, Edward L. Glaeser, and César A. Hidalgo
NBER Working Paper No. 21620
October 2015,Revised October 2015
JEL No. C8,R20,Y10

ABSTRACT

Urban change involves transformations in the physical appearance and the social composition of neighborhoods. Yet, the relationship between the physical and social components of urban change is not well understood due to the lack of comprehensive measures of neighborhood appearance. Here, we introduce a computer vision method to quantify change in physical appearance of streetscapes and generate a dataset of physical change for five large American cities. We combine this dataset with socioeconomic indicators to explore whether demographic and economic changes precede, follow, or co-occur with changes in physical appearance. We find that the strongest predictors of improvement in a neighborhood's physical appearance are population density and share of college-educated adults. Other socioeconomic characteristics, like median income, share of vacant homes, and monthly rent, do not predict improvement in physical appearance. We also find that neighborhood appearances converge to the initial appearances of bordering areas, supporting the Burgess "invasion" theory. In addition, physical appearance is more likely to improve in neighborhoods proximal to the central business district. Finally, we find modest support for "tipping" and "filtering" theories of urban change.

Nikhil Naik
Massachusetts Institute of Technology
75 Amherst Street, E14-374C
Cambridge MA 02139
naik@mit.edu

Scott Duke Kominers
Society of Fellows
Harvard University
Soldiers Field
Boston, MA 02163
kominers@fas.harvard.edu

Ramesh Raskar
Massachusetts Institute of Technology
75 Amherst Street, E14-374K-A
Cambridge MA 02139
raskar@mit.edu

Edward L. Glaeser
Department of Economics
315A Littauer Center
Harvard University
Cambridge, MA 02138
and NBER
eglaeser@harvard.edu

César A. Hidalgo
Massachusetts Institute of Technology
75 Amherst Street, E15-392
Cambridge MA 02139
hidalgo@mit.edu

1 Introduction

How does a city's physical appearance impact its social and economic outcomes? How does a city's social and economic structure shape its physical appearance? The co-evolution between the physical appearance and social composition of cities has interested scholars for centuries. In fact, most architectural and urban planning movements subscribe, at least implicitly, to a theory of the co-evolution between a city's physical appearance and its social composition. As Wilson (1994) wrote, when the City Beautiful movement emerged a century ago, its scholars "trumpeted the ameliorative power of beauty, [. . .] stating their belief in its capacity to shape human thought and behavior." Nevertheless, the relationship between the physical and socioeconomic components of urban change is not well understood due to the lack of comprehensive measures of physical appearance.

We overcome these limitations by using computer vision methods that quantify the physical appearance of streetscapes using images. Specifically, we use "Streetscore", a measure of perceived safety of streetscapes obtained from computer vision methods (Naik et al., 2014) trained on hundreds of thousands of ratings of Google Street View images (Salesses et al., 2013). Streetscore provides a decisive advantage over manually curated or crowdsourced datasets. As a computer algorithm, it has the throughput required to evaluate millions of street blocks to identify physical changes in the urban environment at an unprecedented scale.

In this study, we combine Streetscores with demographic and economic data largely taken from the U.S. Census and American Community Survey. Since neighborhood changes may largely reflect migration, we cannot directly test architectural ideas about how neighborhood appearance impacts the life of specific individuals. Rather, we examine how neighborhood appearance relates to aggregate social outcomes at the neighborhood level. We explore whether social and economic changes precede, follow, or co-occur with changes in the physical appearance of neighborhoods. Additionally, we test three explicit theories of urban change: invasion, tipping, and filtering.

First, we examine whether *ex ante* social characteristics (typically measured in the 2000 U.S. Census) predict changes in the physical appearance of neighborhoods (as measured by changes in Streetscore). We find that the social factors that best predict higher perceived safety in 2007—density and education—are also the best predictors of increases in perceived safety between 2007 and 2014. These relationships hold regardless of whether we control for the 2007 Streetscore or other variables. High housing prices and rents also predict upgrading in Streetscore, perhaps because it is more sensible to invest in the physical environment where land is intrinsically more valuable. Surprisingly, we find that other variables, such as income, race, and poverty, have little or no predictive power in our context.

Second, we examine whether *ex ante* physical characteristics predict changes in neighborhoods' social characteristics. We find limited evidence for this hypothesis. Unfortunately, our measures of social change are limited to the small and overlapping samples from the American Community Survey instead of the U.S. Census. Streetscores in 2007 predicts increases in density, income, and housing prices, as well as decreases in poverty rates and in the share of the African American and Hispanic population. Oddly, higher Streetscores are associated with less growth in college education, perhaps because of impecunious college graduates headed into scruffier neighborhoods.

Third, we study the correlation between changes in Streetscores and changes in social characteristics. We find these interactions tend to be statistically significant, but relatively modest in magnitude.

Finally, we test the invasion, tipping, and filtering hypotheses of urban change. The invasion hypothesis (Burgess, 1925) is the idea that improvements in a neighborhood will spillover to adjacent areas. Like Kolko (2010), we find that improvements are much larger in areas that are initially surrounded by perceptually safer or better-educated neighborhoods. Yet our findings also support the presence of spatial spillovers, as we observe that physical changes in the urban environment tend to cluster spatially, an empirical fact that is also supportive of the invasion hypothesis.

The tipping hypothesis developed by Schelling (1969) and Grodzins (1957) suggests that neighborhoods in bad physical condition will get progressively worse, while nicer areas will get better. Our results show weak evidence for this hypothesis, as we find a positive association between a neighborhood's initial Streetscore and the subsequent Streetscore growth. Yet, we do not find that initially threatening neighborhoods get worse—only that they improve less.

The filtering hypothesis (Margolis, 1982), on the other hand, suggests cycles in which neighborhoods gradually decay until they get upgraded. We also find weak support for this idea, finding that neighborhoods with newer housing stock score higher than neighborhoods built in the 1950s. However, we cannot rule out the possibility that this finding is also reflective of differences in the perception of various architectural styles, since neighborhoods built before 1939 (prior to the widespread adoption of modernist architecture in the U.S.) also score highly.

In summary, we find that population density and share of college-educated adults are the strongest predictors of subsequent improvement in a neighborhood's physical appearance. We also find strong evidence in support of the invasion theory, but only weak evidence in support of tipping and filtering.

2 The Physical City and the Social City

Much has been written on the relationship between the physical environment and societal behavior. Winston Churchill's famous quotation "we shape our buildings, thereafter they shape us," is a particularly pithy statement of the ancient architectural belief that human life is shaped by the built environment. During the Renaissance, Alberti wrote that "beauty will have such an effect, even upon an enraged enemy, that it will disarm his anger" (Lefaivre and Tzonis, 2004), suggesting that violent crime will be lower in more aesthetically attractive areas. Alberti's aesthetic determinism anticipated the City Beautiful movement, where Charles Mulford Robinson (1899) argued that beautiful cities produce "better citizens, because [they are] better instructed, more artistic and filled with civic pride."

While architectural theorists generally trumpet the virtues of beautiful spaces, they often disagree heartily about what makes a space beautiful. Ebenezer Howard (1902) argued for the combination of town and country in lower density Garden Cities. On the other hand, The Radiant City of Charles Jeanneret (a.k.a. Le Corbusier) abstracted cities as "machines for living". He pushed for the development of low-cost, single-use modernist spaces that would optimize economic and social efficiency. Frank Lloyd Wright (1932) responded to Le Corbusier by proposing a semi-suburban Broadacre City which would provide "more light, more freedom of movement, and more general spatial freedom in the establishment of what we call civilization."

In a different strand of literature we find the empirical work on urban change. There is much work on the statistical properties of urban social change, dating back to Weber (1899), Du Bois (1899), and Drake and Cayton (1970). There is also a recent, albeit smaller, literature on the statistical properties of neighborhood appearance. Quantitative attempts to measure the appearance of neighborhoods began with pioneering urban planning scholars, such as Amos Rapoport (1969) and Kevin Lynch (1960). These efforts have continued with more recent work by social scientists (e.g., Sampson et al. (1997); Hwang and Sampson (2014)).

Traditional efforts to quantify neighborhood appearance used visual perception surveys, in which people were asked to rate or compare images. These methods created the first perceptual and evaluative maps of cities at relatively low resolutions (Lynch, 1960; Milgram, 1976; Nasar, 1998). However, the manual nature of traditional data collection processes meant that they could not be deployed over large territories or different time periods.

More recently, the availability of online images and advances in computer vision and machine learning have made it possible to quantify the visual appearance of streetscapes at unprecedented scales. Crowdsourcing increases the throughput of visual perception surveys from hundreds of images to thousands (Salesses et al., 2013). In turn, these crowdsourced studies provide training data for computer vision algorithms that simulate the assessments of humans and score millions of images (Naik et al., 2014). Scores derived from computer vision make it possible to examine which urban designs are more conducive to the development of safer-looking urban environments (Harvey et al., 2015) and to identify the streetscape features that explain perceived safety (Porzi et al., 2015).

The economics and sociology literature on urban change has typically reflected contemporaneous urban events. Weber (1899) and Burgess (1925), for instance, wrote during an era of unbridled urban growth and focused primarily on urban expansion. After World War II, the changing racial character of America's neighborhoods seemed to be the central fact of urban life, and so Grodzins (1957), Taeuber and Taeuber (1969), and Schelling (1969) wrote about racial change and tipping points. When urban decay set in, Hoover and Vernon (1959) and Muth (1969) discussed neighborhood filtering, and Wilson (1978) described the later exodus of prosperous African-Americans from urban ghettos.

Just as many architects often argue that appearance shapes society, many social scientists argue that society determines the changes in the physical city. For example, Burgess (1925) presented a concentric-ring theory of cities (a 20th-century variant on von Thünen and Hall (1966)). This model describes a city in which a central business district (CBD) is surrounded first by a transition zone (containing industry and the poor), then by a working man's residential zone, then by a more prosperous zone, and finally by a zone for commuters. The Burgess model of neighborhood change is one of invasion, in which inner zones expand outward as the overall growth of the city continues. While the driving force is the overall city expansion, the impact on individual neighborhoods can be physical. For instance, previously attractive residential zones can get occupied by industry and lower-income housing.

The Burgess (1925) invasion hypothesis predicts that proximity to the CBD could well be associated with declining perceived safety, as transition zones expand out to residential areas. A 21st-century variant of Burgess, especially for New York and Boston, would note the emergence of prosperous zones right near the city center (Glaeser et al., 2008). The prediction here is that, in thriving cities, the zone near the city center grows and invades its neighbors. In either case, we can test the invasion theory by looking for neighborhood improvements either close to the CBD or in areas that are less attractive, but proximal to more attractive areas.

Another theory about the evolution of a city's physical environment is the famous tipping model of Grodzins (1957) and Schelling (1969). Applied to appearance, their idea is that over time, attractive places continue to improve, while unattractive places tend to get worse. In filtering models (Muth, 1969; Bond and Coulson, 1989), housing decays over time. Consequently, housing is built initially for the prosperous, but as it decays it is sold or rented to the poor. More recently, Resseger (2013) has also shown the power of physical space to shape social characteristics, by documenting that neighborhoods contain fewer minorities when regulations prevent the construction of multi-family housing neighborhoods.

Last but not least, probably the most famous hypothesis connecting urban perception and human behavior is the Broken Windows Theory (BWT) of Wilson and Kelling (1982). The BWT is now used to describe two ideas: (i) urban disorder could trigger disorderly behavior, and (ii) policing minor offenses is an effective way of deterring larger offenses. While BWT is a topic of wide interest that has been explored using experimental (Keizer et al., 2008) and observational methods (Salesses et al., 2013), it is also a topic that goes beyond the scope of our current effort.

Our work focuses on linking the changes in neighborhood appearance, derived from Streetscore, with neighborhood demographic and economic data. We test broad hypotheses about whether the social attributes of a neighborhood predict changes in appearance, whether appearance predicts social changes, and whether both appearance and society change together. In addition, we test the theories of invasion, tipping and filtering. Therefore, our work builds on both the literature on neighborhood appearance and urban change.

Next, we describe our data sources and methods to quantify neighborhood appearance.

3 Data and Methods

In this section, we describe our method to measure changes in neighborhood appearance using computer vision by processing time-series image data from Google Street View.

We query the Google Street View Image API to obtain $360°$ panorama images captured by the Street View vehicles. Each panorama is associated with a unique identifier ("panoid"), latitude, longitude, and time stamp. We identify a section of the panorama by specifying the heading and pitch of the camera relative to the Street View vehicle. We obtained a total of 1,645,760 panoids for street blocks in Baltimore, Boston, Detroit, New York, and Washington D.C. for 2007 and 2014. We call the 2007 panel the "before" panel and the 2014 panel the "after" panel. For the street blocks that lack images for either 2007 or 2014, we completed the "before" and "after" panels using images from the closest years for which data was available. As a result, 5% of the images in the "before" panel are from either 2008 or 2009. Similarly, 12% of the images in the "after" panel are from 2013. Finally, we match images from the "before" and "after" panels by using their geographical locations (i.e. latitude and longitude) and by choosing the same heading and pitch.

This process gave us images that show the same place, from the same point of view, at different times (see Figure 1-(b-d)).

We preprocess the visual data to identify whether one or both of the images in "before" and "after" pairs are over-exposed, blurred, or occluded images and discard such image pairs. Thus, we reduce our data to a set containing only valid "before" and "after" image pairs (for details, see Supplemental Material).

Next, we calculate the perception of safety for each image using a variant of the Streetscore algorithm (Naik et al., 2014) trained on a crowdsourced study of people's perception of safety based on 4,109 images and 208,738 user ratings (Salesses et al., 2013). The perception of safety correlates highly with the perception of wealth (Salesses et al., 2013). We choose to use safety scores because our dataset on perception of safety has a larger number of user ratings than our data set on perception of wealth. As a result, the dataset on perception of safety is more robust for training a machine learning algorithm. Our variant of Streetscore evaluates streetscapes by first segmenting images (Hoiem et al., 2008) into four geometric classes—ground (which contains streets and sidewalks), buildings, trees, and sky (Figure 1). Then Streetscore characterizes each class using two popular computer vision features—GIST (Oliva and Torralba, 2001) and Textons (Malik et al., 2001). Next, Streetscore predicts the crowdsourced rating of an image using support vector machine. To avoid seasonal effects, Streetscore quantifies infrastructural changes alone, that is, changes in ground and buildings. Finally, we compute the changes in Streetscore between images in the "before" and "after" panel from the same location, to obtain the Urban Change Coefficient (UCC)(for details, see Supplemental Material). We aggregate the Streetscores from 2007 and 2014, as well as the UCC between these periods at the census tract level for the five cities in our study.

To relate the Streetscore indicators of neighborhood appearance to socioeconomic data at the census tract level, we obtain tract characteristic data from both the 2000 and 2010 U.S. Census, adjusted to the 2010 census tract boundaries. In addition, we use five-year data from the American Community Survey between years 2006-2010 and 2009-2013. To measure the share of vacant addresses in a census tract, we use the HUD USPS dataset (U.S. Department of Housing and Urban Development, 2015).

In the next section, we describe our findings on relating neighborhood appearance to socioeconomic characteristics.

4 Results

4.1 Basic Characteristics of Physical Urban Change

Here we present the descriptive statistics of Streetscores from the five cities considered in our study, along with maps of spatial distribution of urban changes.

Table 1 provides means and standard deviations for the Streetscore data for 2007 and 2014, along with other summary statistics. The data shows a significant improvement in Streetscores in the five sample cities between 2007 and 2014. The mean of the measure was 7.75 in 2007 and 9.14 in 2014. This 1.3 unit change is almost exactly one half of a standard deviation of the measure in the cross-section in 2007. The standard deviation also increased by 0.78 between the two periods, implying that both the mean and the variance of street appearance has increased.

Next, we present maps illustrating the instances of major urban change seen in the cities of Baltimore, Boston, New York, and Washington D.C. (Figure 2). We measure urban change by defining the Urban Change Coefficient (UCC) as the difference between the Streetscores of the "before" and "after" images for each pair. Positive UCCs are indicative of urban change in which the "after" image (2014) is perceived to be safer than the "before" image (2007). The maps in Figure 2 show instances of large urban change, that is, UCCs representing changes of at least four standard deviations. Our UCC maps show that urban change has a tendency to cluster, meaning that places that experience urban change tend to be close to other places experiencing urban change.

Next, we move away from descriptive statistics and will look at the co-evolution between the physical appearance and social composition of neighborhoods.

4.2 The Physical City and the Social City

In this analysis, we first explore the cross-sectional correlation between a city's physical appearance and social composition. Then, we examine whether a city's physical appearance precedes, follows, or co-occurs with social change.

Our unit of observation is the census tract. We use tract characteristic data from the 2000 U.S. Census, the 2010 U.S. Census as well as 5-year data from the American Community Survey between years 2006-2010 and 2009-2013. In addition, we use the share of vacant addresses in a census tract as reported by the HUD USPS dataset (U.S. Department of Housing and Urban Development, 2015), averaging the data from all quarters of 2007 and 2014.

In all regressions, we have corrected for spatial correlation in standard errors following Conley

(2008) and using STATA routines developed by Hsiang (2010). These corrections are meant to address the clustering of both neighborhood attractiveness and the social characteristics discussed above.

4.2.1 Cross-sectional Correlations between a City's Physical Appearance and Social Composition

Table 2 shows the correlation between Streetscore and neighborhood social characteristics in 2007. Each entry illustrates a coefficient from a separate regression. The first column shows the coefficients and standard errors from a univariate regression. The second column shows the coefficients and standard errors from a regression that also includes indicator variables for the different cities. We have also run multivariate regressions, but we omitted these results to save space.

The first two rows of Table 2 show the most robust cross-sectional correlates: population density and education. The first column shows that without the city controls, as the log of density increases, Streetscores improve by 1.23, which represents about one-half of a standard deviation of Streetscore (2.6). When we control for city-level indicator variables, the coefficient drops to .72, but remains extremely significant. This result is robust to including all of the other controls in the table.

Since Streetscores are roughly linear in log density, the overall relationship is concave. Perceived safety rises with density but the effect levels off. This non-linearity helps make sense of the result. Urban streetscapes are perceived as safe when there are people who live there. Commercial or vacant spaces are seen as threatening. This fact lends support to the famous Jacobs (1961) claim about how safety is increased by "eyes on the street." That said, we cannot claim that dense urban spaces are seen as safer than low-density suburban or rural areas, as we do not have such low-density spaces in the sample. These regressions imply only that urban spaces with high population densities are perceived as being safer than urban spaces with low population densities in five (generally dense) eastern U.S. cities.

The second robust correlate of perceived safety is education. As the share of the population with a college degree increases by 20 percent (one standard deviation), perceived safety rises by .64 without city controls and by .52 with city controls. Again, the result is quite robust. In this case, we suspect that the relationship reflects the tendency of educated people to be willing to pay for safer neighborhoods rather than the ability of educated residents to make a neighborhood feel safe.

There is no robust relationship between median income and perceived safety. The coefficient in the first column is 1.3, which is considerable and significant when the regression is run using ordinary least squares. The correction for spatial correlation of the errors, however, causes the estimated standard errors to rise significantly and leads the coefficient to become statistically insignificant (albeit still fairly large). With the addition of city-level indicator variables, the magnitude becomes much smaller. As in so many other areas (see, e.g., Glaeser and Saiz (2004)), education is a far stronger correlate of urban strength than income.

The fourth and fifth rows show the relationship between Streetscore and housing costs. The fourth row shows the connection with monthly rent. As rents increase by one standard deviation (245), Streetscore increases by 1, without city controls, and by .5 with city controls. Without city controls, a one standard deviation increase in housing prices (1.5) increases Streetscore by almost 5. The coefficient falls dramatically when we control for city, because there are enormous gaps in prices between New York and either Detroit or Baltimore. Within cities, a one standard deviation increase in prices is associated with an increase in Streetscore of slightly more than 2.

While the price regressions have been run with Streetscore as the dependent variable for consistency with the other analyses in Table 2, it may be more appropriate to interpret these results as reflecting the willingness to pay for a visually safer or more attractive environment. If regressions are re-run with city level controls and with the housing costs as the dependent variable, we find that a one standard deviation increase in Streetscore is associated with an increase in housing prices of .4 log points and increased monthly rent of $66. These increases are significant and suggest the value in pursuing serious hedonic regressions using Streetscore, but interpreting our result here is difficult partially because the rent data precedes the Streetscore data by seven years. (Other reasons to be wary include the lack of other housing or neighborhood controls, the self-reported nature of the housing prices, and the issues surrounding rent control in New York.)

The next rows look at race and poverty. The sixth row of Table 2 shows that there is no significant relationship between poverty and Streetscore, which is perhaps unsurprising given that income and Streetscore are not correlated. There is a modest negative relationship between Streetscore and percent African-American. The relationship between Streetscore and percent Hispanic is positive, without the city-level controls, and essentially zero with these controls.

The final row shows the relationship between vacancy and Streetscore. Without city-level fixed effects, there is a negative relationship between vacancy rates and Streetscore. The effect switches sign and becomes stronger in the second column, which may reflect the fact that vacancy rates are much higher for higher-density rental properties than for lower-density, single-family detached homes.

4.2.2 Changes in a City's Physical and Social Composition

We now turn to relationships between changes in a city's physical appearance and social composition. Unfortunately, we do not have exogenous variables that would enable us to tease out the chains of causality. We can, however, examine whether social factors predict physical change, or whether physical factors predict social change. We will also look at the correlation between the two types of change.

Do social factors predict physical change? Table 3 examines the correlation between initial social characteristics, measured with the 2000 U.S. Census and changes in Streetscore between 2007 and 2014. As in Table 2, each entry in Table 3 is a separate coefficient from a different regression. The first column shows the results for density, and the first coefficient means that as the log of density increases by 1, the growth in Streetscore increase by .2 points, which is about one-fourth of a standard deviation of the UCC (.77). The density coefficient is large, but as the second column shows, it primarily reflects differences across cities. When we control for city fixed effects, the coefficient drops to .08, which is still statistically significant and reasonably meaningful. The estimated impact of log density on the change in Streetscore over a seven year period is about one-tenth of the impact of log density on the level of Streetscore in 2007.

The density coefficient drops slightly when we control for the initial Streetscore level in the third column. In the fourth column, we drop New York from the sample and the coefficient remains close in value but loses statistical significance. Density does appear to predict growth in Streetscore over the sample period, but the relationship is far weaker than the connection between density and the level of Streetscore in 2007.

The next row of Table 3 examines education. In this case, the city fixed effects have less of an impact. The coefficient is .008 without city fixed effects and .007 with city fixed effects. All results in this row are quite statistically significant. The coefficient drops only slightly when we control for initial Streetscore. When we exclude New York in the fourth column, the coefficient falls to .0046, but remains statistically robust. Indeed, the correlation between education and growth in Streetscore is robust to including any other variable that we have in our dataset.

Moreover, the education effects are reasonably large. The coefficients on the seven-year change are about one-fourth of the comparable coefficient on the level of Streetscore in 2007. A one standard deviation increase in share with college degree in 2000 (20 percent) is associated with a .13 increase in UCC, or about one-sixth of a standard deviation. Just as skilled cities have done particularly well over the last 50 years (Glaeser et al., 1995), skilled neighborhoods seem to have disproportionately experienced physical improvements.

By contrast, the relationship between UCC and median income is far weaker. The relationship is statistically significant but disappears when we control for city fixed effects. This change reflects the fact that Streetscore improvement was strongest in New York, which is also the richest city. Income is insignificant in all of the specifications where we control for city-level fixed effects.

The fourth and fifth rows of Table 3 look at housing costs. Both monthly rents and the log of median housing price are significantly correlated with Streetscore changes in three of the four specifications. Both variables lose statistical significance when we drop New York, which is perhaps unsurprising given that New York comprises more than half of our sample.

The coefficient on rents is about one-seventh of the comparable coefficient from the comparable levels regression in Table 2. The coefficient on log housing price is about one fourth of the coefficient on log Housing Price in Table 2. A one log point increase in housing prices is associated with a .35 increase in Streetscore, which is reasonably large. One plausible explanation for this strong relationship is that Housing Price is forward-looking, and buyers in 2000 anticipate which neighborhoods will improve. Yet the correlation with rents—which should not anticipate future gains in neighborhood quality—suggests that forward-looking behavior does not explain the entire story.

The next three regressions look at poverty rates, and race. Poverty is not correlated with UCC. Streetscores increased less in areas where there are more African-Americans, but the effect is modest. A 20 percent increase in the share of an area that is African-American is associated with a .04 decrease in Streetscore growth. The coefficient shrinks and becomes statistically insignificant when we exclude New York. The share of the population that is Hispanic does not correlate with changes in Streetscore across the entire sample. The correlation is positive outside of New York.

The last row examines the relationship between the vacancy rate in 2007 and change in Streetscore. In the first column, we find that lower vacancy rates do predict improvements in Streetscore. Yet this relationship disappears once we control for city fixed effects. The result in the first column seems to be driven primarily by the fact that vacancy rates are lower in New York, and New York experienced the largest gains in Streetscore.

Overall, our results on whether social characteristics predict physical change are mixed. Education is certainly a strong predictor. Housing costs and density also predict changes in Streetscore, at least when New York is included. Yet the other demographic variables are generally not significantly correlated with neighborhood change.

Do physical factors predict social change? We now examine whether physical factors predict social change. In Table 4, we look at the connection between Streetscore in 2007 and changes in social variables. The largest difficulty we face here is in measuring changes in social variables over the sample period. Whereas the 2000 U.S. Census has quite good coverage of the relevant tract-level variables, to measure changes between 2007 and 2014 we must rely on the American Community Survey (ACS), which represents a much smaller sample than the U.S. Census does.

We are forced to use five-year averages over the periods 2006-2010 and 2009-2013. (The latter period is the last period for which ACS data is available.) There are several unfortunate aspects of reliance on the ACS. The middle of our second period is 2011, which is only three years after the middle of the earlier period. There is some overlap in the two time periods, but we lose approximately one-quarter of the samples if we use the five-year averages over 2005-2009 as the first period, due to census tract boundary changes. While we are uncomfortable with these issues, we know of no other way to get tract-level changes in social variables during the period of our Streetscore data.

We regress the change in the social variables on the 2007 value of Streetscore. As in Table 2, each entry represents a coefficient from a separate regression. The top left entry shows that a one-standard deviation increase in the initial Streetscore value (2.6) is associated with a .02 log point increase in population density over our time period. In the second column, we see that this coefficient drops in half when we include city level fixed effects. The coefficient increases again when we control for the 2000 density level of the variable in the third column and when we drop the New York data in column (4).

To us, this effect seems almost implausibly large. Taking the third column as our preferred specification, we find that a one standard deviation increase in Streetscore is associated with an over 3 percentage point increase in density over this short time horizon. It does appear that people are moving into areas that are seen as being safer. The results are also quite statistically robust. The next row shows, somewhat surprisingly, that higher Streetscores are associated with less growth in the share that is college educated. In the first two columns, the estimated coefficient is -1, meaning that an increase in Streetscore of 2 is associated with 2% decrease in the growth of the share that is college educated. In the last two columns, where we control for the share of the population that is college educated in 2000, the coefficient drops to -.6. But even this last coefficient is remarkably large for such a short time period.

One possible explanation for this relationship is that the observed effect is showing gentrification in action: better educated people move into areas that are initially less attractive, causing large jumps in the measured skill levels. That hypothesis would seem to suggest that there will be a large positive association between changes in Streetscore and changes in education; we test this later. Yet it is also possible that the young people who are moving into lower-quality areas are relatively transient individuals, such as recent graduates or students. In the later case, we may not expect to see much upgrading of the physical streetscape.

The third row shows the relationship between initial Streetscore and income growth. In all four specifications, a 2 point increase in Streetscore is associated with about a .01 log point increase in incomes. Again, this effect seems relatively large to us since the time period is so small. It does seem to be that incomes are rising faster in places that are initially more attractive.

It is somewhat unusual to see such divergent trends for income and education, which usually go together. One possible explanation is that while changes in education seem more likely to be explained by migration, changes in income seem more likely to reflect changes in the fortunes of longer-term residents. According to this view, lower initial Streetscores are associated with an influx of well-educated gentrifiers, while high initial Streetscores are associated with households whose fortunes were more resilient during the Great Recession.

The fourth and fifth rows of Table 4 show the correlation between Streetscore and changes in housing costs. The correlation with rents is small and statistically insignificant overall, although it becomes larger and significant when we exclude New York. The correlation with changes in housing prices is more robust. A one standard deviation increase in Streetscore in 2007 is associated with a growth in housing prices of .006 log points in our preferred specification (the third column). Given the enormous volatility of housing prices over this time period, this change seems relatively modest.

The next three rows of Table 4 show the relationship between changes in Streetscore and changes in poverty and race. In all three specifications with city fixed effects, a one standard deviation increase in Streetscore in 2007 is associated with a -.4 percentage point reduction in the poverty rate. It is hard to discern whether this represents the movement of wealthier people into nicer areas, or whether it reflects people in nicer areas experiencing less economic distress over this period.

The results on race, naturally, can reflect only migration. In our preferred specification, a one standard deviation increase in Streetscore is associated with a reduction in percent African-American of .18 percent. The comparable regressions in the last row shows that a one standard deviation increase in Streetscore is associated with a .26 percentage point reduction in the share of Hispanics. The effects are not huge, but they do seem to suggest that whites have moved disproportionately into areas that appear safer.

The last row shows the relationship between changes in vacancy rates and the initial Streetscore value. The first column shows that a one standard deviation increase in the initial Streetscore is associated with a quite significant 1% reduction in the vacancy rate. This coefficient loses three quarters of its magnitude when we control for city level fixed effects, but remains statistically significant. The coefficient becomes insignificant in the third column, but then jumps back up in significance and magnitude when we exclude New York. It appears that Streetscore is a solid predictor of changes in the vacancy rate outside of New York, but not in New York, where the vacancy rates tend to be extremely low.

While we recognize the limitations of our ACS sample, our results do seem to suggest that Streetscore predicts changes in the social environment. These results are typically robust to other controls, so they do at least suggest that the social characteristics of a neighborhood respond to the physical state of the area.

What are the correlations between physical and social change? Table 5 turns to the correlations in physical and social trends. The top three columns show the relationship between 2007-2014 Streetscore changes and 2000-2010 census changes. The remainder relates the same Streetscore changes to the changes in ACS variables used in Table 4. The advantage of the census variables is that they are more precisely measured. The disadvantage is that the time period matches less well to the Streetscore change (and we have only three usable variables). Again, each entry shows the result of a separate regression. We have regressed Streetscore changes on changes in the social variables, but no causality is implied by this choice. These regressions merely test whether there is a parallel trend in the two variables.

The top row shows the correlation between changes in Streetscore and changes in density between 2000 and 2010. The coefficients range from .03 to .07. In our preferred specification, a density increase of .2 is associated with a .008 increase in Streetscore. This effect is modest, but it is statistically significant. The fourth row reproduces this relationship using the ACS and finds that coefficients drop by two thirds but remains marginally significant in our preferred specification (the third column).

The next two rows show the results for changes in the shares of the population that are African-American and Hispanic. Again the results are generally significant, with more growth in either minority group being associated with less growth in perceived safety. The effects seem modest to us: a 10 percent point increase in the share of the population that is African-American in associated with a .07 point decrease in Streetscore, which is slightly less than one-tenth of standard deviation. The coefficient on change in Hispanic share is somewhat stronger. A ten-percentage point increase in the share of Hispanics is associated with a .13 decrease in Streetscore.

The remaining rows all use data from the ACS. The change in the share with a college degree is negatively associated with changes in Streetscore. Again, this is somewhat surprising and perhaps explained by young, more transient populations moving into areas with initially lower Streetscores.

Changes in income are positively associated with changes in Streetscore. Changes in poverty are negatively associated with changes in Streetscore. There is no correlation with changes in housing prices or rents. It is hard to know if these weak results reflect a genuine lack of correlation or whether this reflects the weaknesses of our overlapping American community survey sample. Changes in the vacancy rate (obtained from the HUD USPS dataset) are strongly associated with changes in Streetscore across the entire sample if we do not control for city-level fixed effects, but become insignificant once we include these basic controls.

Overall, though, we do see trends in changes going together, especially when we use the considerably more accurate census data. These results support the view that there are interactions between the physical city and the social city.

4.3 Three Theories of Neighborhood Change

We now test three famous theories of neighborhood change: invasion, tipping, and filtering.

4.3.1 Invasion

We first test the Burgess (1925) invasions model. The natural test for invasion is to regress changes on the characteristics of bordering neighborhoods and proximity to downtown. We begin by controlling for proximity to the Central Business District (CBD) in each city, with the locations of the CBDs based on coding of Holian and Kahn (2012). As in Table 2, each entry in Table 6 represents a coefficient estimated in a separate regression.

The left column shows the results for all the cities with no controls whatsoever. In this case, the coefficient is essentially 0. The next column includes city indicator variables and the estimated coefficient falls to a statistically significant -.07. This switch reflects the fact that New York had large Streetscore gains and average distance to CBD is much higher there than in smaller cities like Baltimore and Boston.

This means that as the distance to the CBD increases by one mile, the expected growth in the Streetscore measure falls by .07 points. The third column shows that the coefficient changes little if we also control for initial education, log of density, and Streetscore. The fourth column shows that the coefficient becomes slightly more negative when we drop New York from the sample. (We drop New York because it represents more than 50% of our data, and it is important to make sure that the results are not simply reflecting phenomena of a single large city.)

The original invasion hypothesis postulated a process under which urban slums would gradually make their ways out from the center to nearby suburbs. The current pattern is instead one in which the central city sees particularly large upgrades in perceived street safety. Our results are similar to those of Kolko (2010), who also found that gentrification was closely correlated with distance to the city center (although this may be a characteristic of cities in developed countries, and not of cities in general). One interpretation is that we are currently witnessing the reversal of the process described by Park and Burgess. City centers, which always had a strong fundamental asset—proximity to jobs—are reverting to that fundamental.

While the data does not suggest decay emanating out from a center, the core idea of the invasions hypothesis—that neighborhoods spill over into each other—is readily vindicated by the data. The next row of Table 6 regresses UCC on the average Streetscore in surrounding areas. The coefficient without city indicators is .14, but in all three of the other specifications, the coefficient is around .09. Notably, the coefficient on neighboring scores is more than double the impact of the neighborhood's own score, implying that almost one-tenth of the Streetscore difference between a neighborhood and its neighbors is eliminated over a seven year period. Since most of the movement over the sample period is positive, the regression should be interpreted as meaning that growth is faster in areas with more attractive neighbors. This strong convergence is exactly the prediction of the invasion theory.

The third row of Table 6 examines the effect of adjacent density. As adjacent density is highly correlated with adjacent Streetscore, it is not surprising to see that there is also a robust correlation here, although the connection is not as strong as with adjacent Streetscores. The fourth row examines the average house price in the adjacent areas. In this case, the correlation is significant without controls, but proves not to be robust to controlling for the initial area attributes.

Finally, the fifth row looks at average share of the population with college degrees in adjacent areas. The relationship is positive and robust. As the share increases by 20%, Streetscore increases by .12 points. This again corroborates the results of Kolko (2010), who found that gentrification is faster in areas with more educated neighbors. These findings point to a process of neighborhood convergence, which is, in a sense, at the heart of the invasion hypothesis.

4.3.2 Tipping

In tipping models, neighborhoods either spiral downwards or climb upwards. In extreme versions of the model, there is a single threshold below which a neighborhood moves downward and above which a neighborhood moves upward.

To test the tipping theory, we look at the relationship between a tract's 2007 Streetscore and that tract's change in Streetscore between 2007 and 2014. Our first test is simply to regress the change in Streetscore between 2007 and 2014 on the initial Streetscore level. If the tipping theory is correct, then this coefficient should be positive. We run this regression with city indicator controls and there are 2514 observations. We estimate:

$$UCC = \frac{1.22}{0.942} + \frac{0.043}{0.011} \cdot Streetscore2007 \tag{1}$$

The coefficient on 2007 Streetscore is certainly positive. For each extra point of Streetscore in 2007, Streetscore growth is .04 points higher over the next seven years. Yet it is not clear whether this represents tipping or a pattern where visually safer areas are being upgraded first and faster.

To look for a tipping pattern, we present the data visually in Figure 3. We group the data into sixteen bins based on the initial value of Streetscore. We determine the average Streetscore growth in each bin and plot the results. The pattern shows the positive relationship shown in the regression, albeit with significantly concavity. If a neighborhood starts with a sufficiently high score, it has little room to move up. Yet while there does seem to an upward spiral, there is little evidence of a downward spiral, since average growth is still positive for the neighborhoods with the lowest initial scores.

We suspect that this lack of downward movement may be particular to the time period under consideration. Despite the Great Recession, 2007-2014 was a relatively good time period for many of America's eastern cities, and this may explain why we do not see declines even for less attractive neighborhoods. Still, the pattern predicted by tipping models, in which upward growth is faster in initially better areas does seem to be supported by the data.

4.3.3 Filtering

Finally, we look at the filtering hypothesis, which suggests the importance of the age of the building stock: areas should gradually decline until they are upgraded. To test the hypothesis that building age shapes streetscape change, we regress Streetscores and changes in Streetscore on the shares of the building stock (as of the year 2000) built during different decades. The data groups together all buildings erected before 1939.

The first regression in Table 7 shows the results for 2007 Streetscore with city indicator variables. The second regression shows results for 2014. The coefficient on the share built between 1990-2000 is strongly positive in both regressions. The coefficients on the share built in the 1980s are also positive, although somewhat weaker. As the filtering hypothesis suggests, newer areas are seen as being more attractive.

The share of the housing stock from the 1970s has a neutral coefficient, but the share of the stock from the 1940s, 1950s and 1960s generally has a negative relationship with Streetscore. The coefficient on the stock from the 1950s is somewhat more negative than the coefficient on the stock from the 1960s, but the coefficient on the stock from the 1940s is slightly less negative than the coefficient on the stock from the 1950s.

But then the coefficient on the stock from the 1940s is positive. Older post-war building seems to have an adverse effect on perceived safety, while pre-war building seems to have a positive impact. One natural explanation for this switch—as suggested by filtering—is that the post-war stock is generally new enough not to have been naturally replaced. The pre-war stock is old enough so that replacement would have occurred unless the buildings are sufficiently attractive. If this explanation is correct, then the replacement creates selection so that the older buildings that survive are particularly nice ones.

The third regression in Table 7 examines the changes in Streetscore. Here we see a slightly positive coefficient for the recent buildings and a negative coefficient for the older post-war buildings. This too is compatible with a filtering view that older buildings are in the process of degenerating. The positive coefficient on pre-war buildings can also be interpreted as meaning that some of these buildings have reached the point where they are being regenerated, either by new buildings or through reinvestment.

The fourth regression controls for initial levels of Streetscore, share of the population with a college degree, and log of density. The results become weaker, but the coefficients on the share of the stock built during the 1980s remains positive and significant and the coefficient on the share of the stock building during the 1940s remains negative and significant. These results again support the general thrust of the filtering hypothesis that older neighborhoods gradually degenerate.

5 Conclusion

We used a new measure of urban perception—Streetscore—to examine the link between the physical city and the social city. We first tested whether social attributes of the neighborhood, including income, education and race, predict changes in perceived safety. Our results were mixed. Education and population density strongly predict changes in the physical environment, but other variables show less correlation. We then tested whether the physical environment predicted changes in the attributes of the humans living in the area. We found strong correlations between changes in community attributes and neighborhoods' initial physical appearance, although we have some concerns about the ACS data. Both population density and incomes rose significantly more in places that were initially perceived as safer. Finally, we documented the correlation between changes in the physical city and changes in the social city.

We also tested three long-standing hypotheses about neighborhood change and found support for all three. We tested the invasion hypothesis by examining whether proximity to the city center or the attributes of adjacent neighborhoods predicts changes in the physical environment—and found that they do. We found the growth in perceived neighborhood safety to be greater in areas that were initially seen as safer. This result lends support to the tipping hypothesis—although we found little evidence for rapid decay in areas with low initial Streetscores. We also tested the filtering hypothesis and found that the age of the building stock was associated with both perceived safety and changes in perceived safety.

To us, our work raises at least as many questions as it answers. Is the correlation between density and perceived safety true more generally? Does the correlation between physical environment and the social neighborhood reflect migration or changes to the people who are permanently living in the area? We hope that future research can address these questions and can continue exploring the links between the physical city and the humans that reside there.

Supplemental Material

Methods

Here we describe the method to compute the Urban Change Coefficient (UCC) from a pair of Street View images of the same location from 2007 (the "before" image) and 2014 (the "after" image). This process consists of two main steps: (1) Removal of unsuitable images and (2) Streetscore prediction. We describe these step below.

Step 1 – Removal of Unsuitable Images

First, we remove over-exposed images. Typically, these images are a result of the sun directly shining into the camera. To identify over-exposed pixels, we covert the image to the CIELAB color space, where L represents lightness and a, b channels represent the color. The color channels are combined as $C = (a, b)^T$. We compute an over-exposed map $\mathcal{M}$, which gives a value between 0 and 1. (Guo et al., 2010). At pixel i,

$$\mathcal{M}_i = \frac{1}{2} \cdot \tanh\left(\delta \cdot \Big((L_i - L_T) + (C_T - ||C_i||_2)\Big)\right),$$

where δ, L_T and C_T are constants with values $1/60, 80,$ and 40, respectively. We obtain the mean $\mathcal{M}_i^{\text{sky}}$ of $\mathcal{M}_i$ over all the pixels that belong to the "sky" geometric class, as predicted by the Geometric Layout algorithm (Hoiem et al., 2008). We discard the image pair if at least one of the images satisfies $\mathcal{M}_i^{\text{sky}} > 0.85$, which indicates over-exposure.

Second, we remove images that are out-of-focus or contain motion blur. To detect blurred images, we compute the Absolute Central Moment (ACMO) of each image. ACMO is a statistical measure that allows a simultaneous optimization of both focus and exposure (Shirvaikar, 2004). If the normalized value of ACMO is less than 0.2, we label the image as blurred. We discard all the image pairs in which at least one of the images is blurred.

Finally, we remove image pairs in which the number of pixels in the image occupied by the "tree" class (Hoiem et al., 2008) change by more than 10% between the "before" or "after" images. We find this an effective way of discarding image pairs where one of the images contains greenery occluding buildings, while the other does not.

Step 2 – Streetscore Prediction

After removing images unsuited for urban change detection, we predict the Streetscore of remaining images using a regression algorithm trained from computer vision features and the crowd-sourced study of Salesses et al. (2013).

Salesses et al. (2013) performed an online study where participants were shown images of streetscapes randomly chosen from New York, Boston, Linz and Salzburg. Participants were asked to choose one of the two images in response to the questions: “Which place looks safer?”, “Which place looks more upper class?” and “Which place looks unique?”. In this study, 7,872 unique participants from 91 countries ranked 4,109 images using 208,738 pairwise comparisons (or “clicks”). Following Naik et al. (2014), we convert the pairwise comparisons to ranked scores using the Microsoft Trueskill Algorithm (Herbrich et al., 2006), which is a Bayesian ranking algorithm. We denote the Trueskill scores for the question “Which place looks safer?” as Streetscores. The images and their Streetscores form the training set for predicting the human perception of new images using machine learning. In this paper, we only use images from New York and Boston for training since we are interested in predicting the human perception of American cities.

To predict streetscores, we begin with a modified version of the Geometric Layout algorithm (Hoiem et al., 2008) to assign pixel-wise semantic labels. We assign pixels to four geometric classes: “Ground”, “Buildings”, “Trees”, and “Sky”. Next, we extract two different image features separately for the pixels of the four geometric classes. First, we use Textons (Malik et al., 2001) to create a finite dictionary of the micro-structures present in each geometric class and encode the textures of objects such as walls and trees. We generate a texton dictionary by convolving the images with a Gaussian filterbank and clustering their responses together. The clusters represent a “word” in the texton vocabulary and every pixel is assigned to the nearest cluster center, creating a texton map. Second, we use the GIST feature descriptor introduced by Oliva and Torralba (2001) on each of the geometric classes. GIST is a global image feature that provides a low dimensional representation of the spatial layout properties of a scene.

After feature extraction, we use ν-Support Vector Regression (ν-SVR) (Schölkopf et al., 2000) to predict the Streetscores of images. Given a set of training images with feature vectors $\mathbf{x}$ and Streetscores $q \in \mathbb{R}$, ν-SVR with a linear kernel generates a weight vector $\mathbf{w}$ and a bias term b under a set of constraints. The two variables $(\mathbf{w}, b)$ are used to predict the Streetscore for a new image with feature vector $\mathbf{x}'$ by evaluating $q' = \mathbf{w} \cdot \mathbf{x}' + b$. We measure the accuracy of our predictor using the Coefficient of Determination (R^2). We obtain $R^2 = 57\%$ over 5-fold cross-validation on the training set.

Armed with the Streetscore predictor, we move to the task of measuring urban change from image pairs. Recall that our predictor is a weight vector trained using image features on top of the four geometric classes: ground, building, trees and sky. Therefore, we are able to compute the contribution of each geometric class to the Streetscore of each image. We choose to discard the contribution of the "trees" and "sky" classes since their scores depend on the season and weather at the time of image capture. Note that the "trees" class contains only large trees (and not landscaping), allowing us to account for changes in the built environment due to landscaping from the "ground" class.

Once we compute the Streetscore for each image in a "before"-"after", we calculate the difference in the Streetscores of the two images to obtain the score for urban change. We call this score the "Urban Change Coefficient (UCC)".

References

Bond, E. W. and N. E. Coulson (1989). Externalities, filtering, and neighborhood change. *Journal of Urban Economics 26*(2), 231–249.

Burgess, E. W. (1925). The growth of the city. In R. E. Park, E. W. Burgess, and R. D. MacKenzie (Eds.), *The City*, Volume 31, pp. 47–62. University of Chicago Press.

Conley, T. G. (2008). Spatial econometrics. In S. N. Durlauf and L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics*. Basingstoke: Palgrave Macmillan.

Drake, S. C. and H. R. Cayton (1970). *Black Metropolis: A Study of Negro Life in a Northern City*. University of Chicago Press.

Du Bois, W. E. B. (1899). *The Philadelphia Negro: A Social Study*. University of Philadelphia Press.

Glaeser, E. L., M. E. Kahn, and J. Rappaport (2008). Why do the poor live in cities? The role of public transportation. *Journal of Urban Economics 63*(1), 1–24.

Glaeser, E. L. and A. Saiz (2004). The rise of the skilled city. *Brookings–Wharton Papers on Urban Affairs*, 47–105.

Glaeser, E. L., J. A. Scheinkman, and A. Shleifer (1995). Economic growth in a cross-section of cities. *Journal of Monetary Economics 36*(1), 117–143.

Grodzins, M. (1957). Metropolitan segregation. *Scientific American 197*(4), 33–41.

Guo, D., Y. Cheng, S. Zhuo, and T. Sim (2010). Correcting over-exposure in photographs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 515–521.

Harvey, C., L. Aultman-Hall, S. E. Hurley, and A. Troy (2015). Effects of skeletal streetscape design on perceived safety. *Landscape and Urban Planning 142*, 18–28.

Herbrich, R., T. Minka, and T. Graepel (2006). Trueskill: A Bayesian skill rating system. In *Advances in Neural Information Processing Systems*, pp. 569–576.

Hoiem, D., A. A. Efros, and M. Hebert (2008). Putting objects in perspective. *International Journal of Computer Vision 80*(1), 3–15.

Holian, M. J. and M. E. Kahn (2012). The impact of center city economic and cultural vibrancy on greenhouse gas emissions from transportation. Mineta Transportation Institute Publications Research Report.

Hoover, E. M. and R. Vernon (1959). *Anatomy of a Metropolis: The Changing Distribution of People and Jobs Within the New York Metropolitan Region*. Doubleday.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the

Caribbean and Central America. *Proceedings of the National Academy of Sciences 107*(35), 15367–15372.

Hwang, J. and R. J. Sampson (2014). Divergent pathways of gentrification racial inequality and the social order of renewal in Chicago neighborhoods. *American Sociological Review 79*(4), 726–751.

Jacobs, J. (1961). *The Death and Life of Great American Cities*. Vintage.

Keizer, K., S. Lindenberg, and L. Steg (2008). The spreading of disorder. *Science 322*(5908), 1681–1685.

Kolko, J. (2010). The determinants of gentrification. Public Policy Institute of California Working Paper.

Lefaivre, L. and A. Tzonis (2004). *The Emergence of Modern Architecture: A Documentary History from 1000 to 1810*. Psychology Press.

Lynch, K. (1960). *The Image of the City*. MIT Press.

Malik, J., S. Belongie, T. Leung, and J. Shi (2001). Contour and texture analysis for image segmentation. *International Journal of Computer Vision 43*(1), 7–27.

Margolis, S. E. (1982). Depreciation of housing: An empirical consideration of the filtering hypothesis. *Review of Economics and Statistics 64*(1), 90–96.

Milgram, S. (1976). Psychological maps of paris. pp. 104–124. Holt, Rinehart, and Winston.

Muth, R. F. (1969). *Cities and Housing: The Spatial Pattern of Urban Residential Land Use*. University of Chicago Press.

Naik, N., J. Philipoom, R. Raskar, and C. A. Hidalgo (2014). Streetscore – predicting the perceived safety of one million streetscapes. In *IEEE Computer Vision and Pattern Recognition*.

Nasar, J. L. (1998). *The Evaluative Image of the City*. Sage Publications.

Oliva, A. and A. Torralba (2001). Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision 42*(3), 145–175.

Porzi, L., B. Samuel, L. Bruno, and R. Elisa (2015). Predicting and understanding urban perception with convolutional neural networks. *ACM Multimedia*.

Rapoport, A. (1969). *House Form and Culture*. Prentice-Hall.

Resseger, M. (2013). The impact of land use regulation on racial segregation: Evidence from Massachusetts zoning borders. Harvard University Working Paper.

Robinson, C. M. (1899). Improvement in city life: Aesthetic progress. *Atlantic Monthly 83*, 771–785.

Salesses, P., K. Schechtner, and C. A. Hidalgo (2013). The collaborative image of the city: Mapping the inequality of urban perception. *PloS one 8*(7), e68400.

Sampson, R. J., S. W. Raudenbush, and F. Earls (1997). Neighborhoods and violent crime: A multilevel study of collective efficacy. *Science 277*(5328), 918–924.

Schelling, T. C. (1969). Models of segregation. *American Economic Review 59*(2), 488–493.

Schölkopf, B., A. J. Smola, R. C. Williamson, and P. L. Bartlett (2000). New support vector algorithms. *Neural Computation 12*(5), 1207–1245.

Shirvaikar, M. V. (2004). An optimal measure for camera focus and exposure. In *Proceedings of the Thirty-Sixth Southeastern Symposium on System Theory*, pp. 472–475.

Taeuber, K. E. and A. F. Taeuber (1969). *Negroes in Cities: Residential Segregation and Neighborhood Change*. Atheneum.

U.S. Department of Housing and Urban Development (2015). HUD aggregated USPS administrative data on address vacancies.

von Thünen, J. H. and P. Hall (1966). *Isolated State*. Pergamon.

Weber, A. F. (1899). *The Growth of Cities in the Nineteenth Century: A Study in Statistics*. Columbia University Press.

Wilson, J. Q. and G. L. Kelling (1982). Broken windows. *Atlantic Monthly 249*(3), 29–38.

Wilson, W. H. (1994). *The City Beautiful Movement*. JHU Press.

Wilson, W. J. (1978). The declining significance of race. *Society 15*(5), 11–11.

Wright, F. L. (1932). *The Disappearing City*. WF Payson.

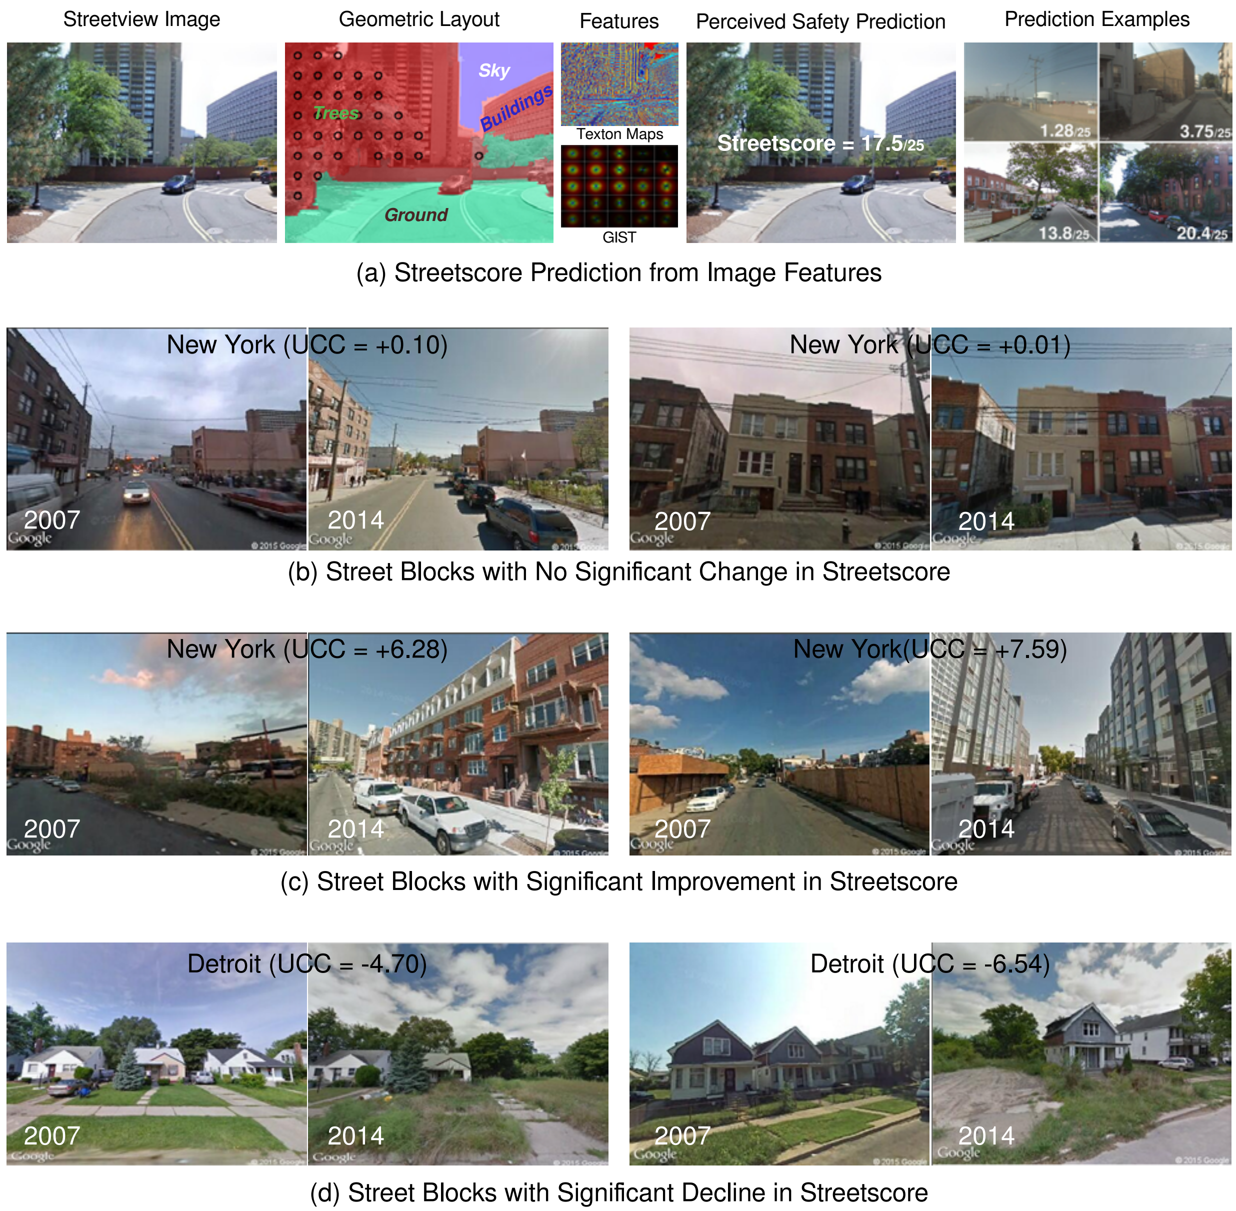

Figure 1: *Computing urban change. (a) We predict the Streetscore of a streetscape by first decomposing Street View images using the "Geometric Layout" algorithm, that splits them into ground, buildings, trees, and sky, and then by using texton map and GIST as image features. (b-d) We calculate the Urban Change Coefficient (UCC) of a street block as the difference between Streetscores of a pair of images captured in 2007 and 2014. (b) UCCs are not affected by seasonal and weather changes. (c) Large positive UCCs are associated with major construction. (d) Large negative UCCs are associated with urban decay.*

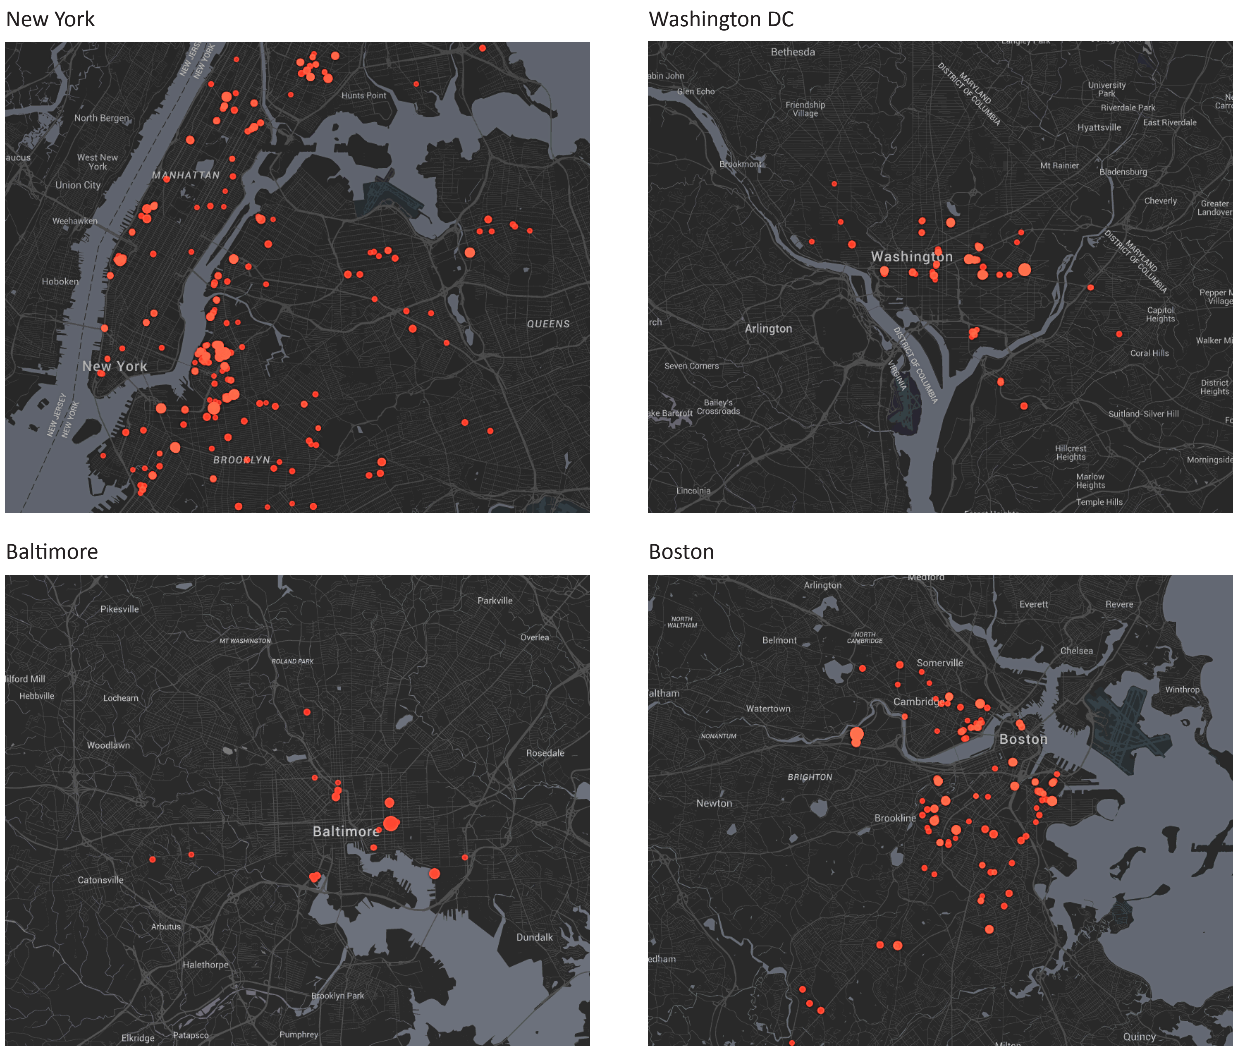

Figure 2: *We present maps of four cities in our dataset, using red circles to indicate locations with UCCs representing changes of at least four standard deviations. These maps show that urban change has a tendency to cluster. (Figure contributed by Mia Petkova.)*

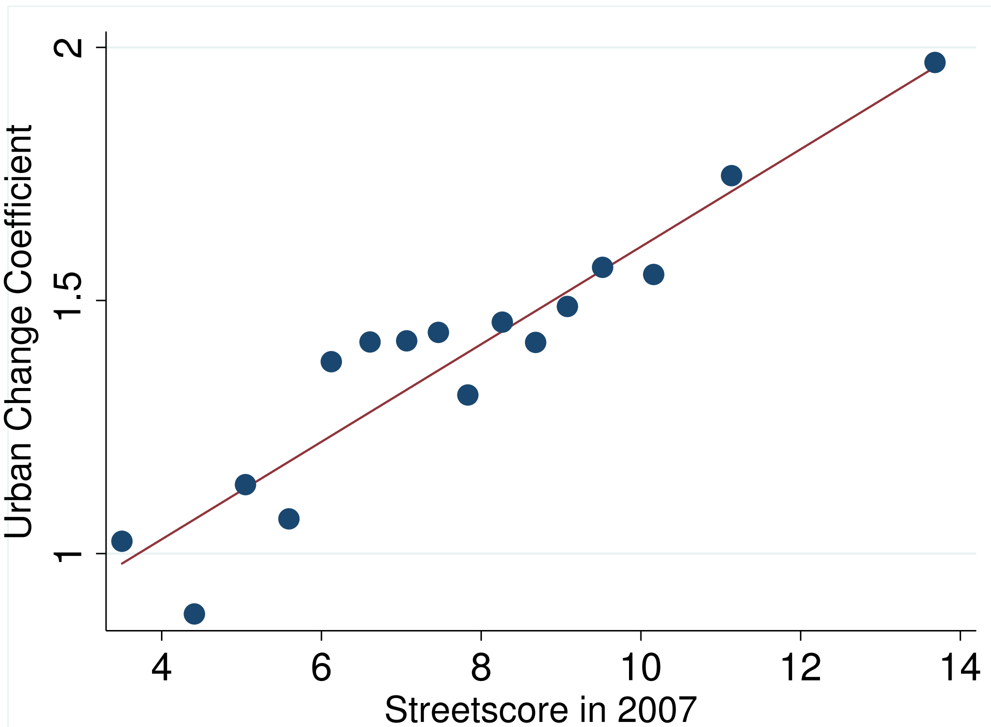

Figure 3: *We test the "tipping" model of neighborhood change. We group the data into sixteen bins based on the initial value of Streetscore and plot the average Streetscore change (UCC) in each bin against the average initial Streetscore.*

Table 1: *Summary Statistics*

Variables	Mean	SD	Min	Max
Streetscore Variables (N = 2514)				
Mean Streetscore of Census Tract, 2007	7.757	2.587	1.681	18.93
Mean Streetscore of Census Tract, 2014	9.146	2.93	2.742	22.12
Mean Urban Change Coefficient of Census Tract	1.39	0.779	-4.076	6.121
Mean Streetsccore of Adjacent Census Tracts, 2007	7.787	2.309	2.548	17.24
Mean Streetscore of Adjacent Census Tracts, 2014	9.192	2.671	3.439	19.4
Mean UCC of Adjacent Census Tracts	1.406	0.568	-1.844	4.647
Socioeconomic Variables from 2000 U.S. Census (N = 2514)				
Log Density	-4.655	1.22	-15.29	-2.48
Mean Log Density of Adjacent Census Tracts	-4.508	0.883	-11.09	-2.73
Share College Education	25.37	21.57	0	100
Mean Share College Education of Adjacent Census Tracts	25.07	19.11	1.149	100
Log Median Income	4.54	0.206	3.884	5.276
Monthly Rent	648.6	253.9	99	2,001
Log Median Housing Price	5.223	0.311	3.938	6
Log Median Housing Price of Adjacent Census Tracts	5.226	0.273	3.833	5.888
Poverty Rate	21.84	13.69	0	100
Share African-American	36.61	37.06	0	100
Share Hispanic	19.22	22.05	0	92.65
Socioeconomic Variables from 2010 U.S. Census (N = 2514)				
Log Density	-4.658	1.25	-14.71	-2.52
Share African-American	35.12	35.84	0	100
Share Hispanic	20.5	21.87	0	100
Socioeconomic Variables from 2006-2010 ACS (N = 2514)				
Log Density	-4.628	1.12	-13.61	-2.45
Share College Education	31.59	23.42	0	100
Log Median Income	4.657	0.224	3.939	5.367
Monthly Rent	980.3	364.5	99	2,001
Log Median Housing Price	5.57	0.331	4	6
Poverty Rate	20.87	14.4	0	100
Share African-American	34.39	36.2	0	100
Share Hispanic	20.22	22.67	0	100
Socioeconomic Variables from 2009-2013 ACS (N = 2514)				
Log Density	-4.619	1.17	-13.99	-2.46
Share College Education	54.72	19.98	0	100
Log Median Income	4.678	0.233	3.99	5.398
Monthly Rent	1,215	377.3	261	2,001
Log Median Housing Price	5.546	0.369	4.167	6
Poverty Rate	21.03	13.5	0	72.11
Share African-American	33.74	35.65	0	100
Share Hispanic	20.41	22.31	0	100

Table 1: *Summary Statistics (Continued)*

Variables	Mean	SD	Min	Max
Share of Housing Stock Built in Various Decades from 2000 U.S. Census (N = 2220)				
Total Housing Stock	1,532	1,007	3	11,522
Share of Housing Stock Built during 1990-2000	5.868	8.744	0	100
Share of Housing Stock Built during 1980-1989	4.043	5.726	0	65.89
Share of Housing Stock Built during 1970-1979	6.941	7.436	0	71.85
Share of Housing Stock Built during 1960-1969	12.3	9.832	0	100
Share of Housing Stock Built during 1950-1959	16.43	10.5	0	74.82
Share of Housing Stock Built during 1940-1949	17.08	9.561	0	81.52
Share of Housing Stock Built before 1940	39.41	20.11	0	100
Vacant Buildings according to the HUD USPS dataset (N = 2514)				
Share Vacant during 2007	3.619	5.528	0	34.74
Share Vacant during 2014	4.267	7.739	0	54.17

Table 2: *Cross-sectional correlations between a city's physical appearance and social composition.*

	Coefficient	
Independent Variable	Without CD	With CD
Log Density	1.152***	0.740***
	(0.082)	(0.095)
Share College Education	0.032***	0.026***
	(0.008)	(0.007)
Log Median Income	1.266	0.145
	(0.797)	(0.706)
Monthly Rent	0.004***	0.002***
	(0.001)	(0.001)
Log Median Housing Price	3.781***	1.344***
	(0.376)	(0.520)
Poverty Rate	0.001	0.011
	(0.009)	(0.008)
Share African-American	-0.024***	-0.005*
	(0.003)	(0.003)
Share Hispanic	0.032***	0.004
	(0.004)	(0.006)
Share Vacant	-0.002	0.013***
	(0.002)	(0.002)

$***p < 0.01, **p < 0.05, *p < 0.1$ *CD – City Dummies.*
All independent variables are from the 2000 U.S. Census. The dependent variable is the mean Streetscore of a census tract in 2007. Standard errors corrected for spatial correlation are in parentheses. Regressions are estimated with a constant that is not reported.

Table 3: *Do social characteristics predict changes in Streetscore?*

Independent Variable	Coefficient Without CD	 With CD	 With CD + SS 07	 With CD + SS 07 (Excl. NY)
Log Density	0.191***	0.077***	0.052***	0.055*
	(0.020)	(0.018)	(0.020)	(0.031)
Share College Education	0.008***	0.007***	0.007***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)
Log Median Income	0.368**	0.140	0.134	0.0134
	(0.150)	(0.132)	(0.117)	(0.177)
Monthly Rent	0.0007***	0.0003***	0.0002**	0.0002
	(0.0001)	(0.0001)	(9.97e-05)	(0.0001)
Log Median Housing Price	0.217***	0.082**	0.071**	0.046
	(0.045)	(0.034)	(0.033)	(0.090)
Poverty Rate	-0.002	0.0003	-0.0001	0.0011
	(0.002)	(0.002)	(0.002)	(0.003)
Share African-American	-0.006***	-0.002**	-0.002**	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Share Hispanic	0.007***	0.001	0.001	0.012***
	(0.001)	(0.001)	(0.001)	(0.003)
Share Vacant	-0.040***	0.007	0.002	-0.008
	(0.004)	(0.006)	(0.005)	(0.006)

$***p < 0.01, **p < 0.05, *p < 0.1$ *CD – City Dummies, SS 07 – Mean Streetscore of Census Tract, 2007.*

All independent variables are from the 2000 U.S. Census. The dependent variable is the mean Streetscore change (UCC) of a census tract between 2007 and 2014. Standard errors corrected for spatial correlation are in parentheses. Regressions are estimated with a constant that is not reported.

Table 4: *Does Streetscore predict changes in social characteristics?*

Dependent Variable	Coefficient Without CD	With CD	With CD + DV 00	With CD + DV 00 (Excl. NY)
Change in	0.009***	0.004**	0.012***	0.016***
Log Density	(0.002)	(0.002)	(0.002)	(0.0037)
Change in	-1.191***	-1.010***	-0.602***	-0.603***
Share College Education	(0.121)	(0.138)	(0.137)	(0.198)
Change in	0.006***	0.005***	0.005***	0.005***
Log Median Income	(0.001)	(0.001)	(0.001)	(0.002)
Change in	-0.060	1.981	2.165	6.255*
Monthly Rent	(1.282)	(1.557)	(1.390)	(3.321)
Change in	0.011***	0.004***	0.003***	0.002
Log Median Housing Price	(0.001)	(0.001)	(0.001)	(0.001)
Change in	-0.104	-0.211***	-0.226***	-0.228
Poverty Rate	(0.066)	(0.076)	(0.072)	(0.167)
Change in	-0.012	-0.063	-0.092*	-0.346***
Share African-American	(0.047)	(0.055)	(0.051)	(0.11)
Change in	-0.117***	-0.140***	-0.132***	-0.099
Share Hispanic	(0.038)	(0.051)	(0.049)	(0.090)
Change in	-0.419***	-0.116**	-0.050	-0.513***
Share Vacant	(0.076)	(0.055)	(0.060)	(0.170)

$***p < 0.01, **p < 0.05, *p < 0.1$ *CD – City Dummies, DV 00 – 2000 Value of Dependent Variable.*

All dependent variables are from the ACS. The independent variable is the mean Streetscore of a census tract in 2007. Standard errors corrected for spatial correlation are in parentheses. Regressions are estimated with a constant that is not reported.

Table 5: *How do social changes relate to physical changes?*

Dependent Variable	Coefficient Without CD	 With CD	 With CD + DV 00	 With CD + DV 00 (Excl. NY)
	Panel A: Regressions using 2000 U.S. Census data			
Change in	0.064***	0.034***	0.040***	0.070***
Log Density	(0.010)	(0.010)	(0.010)	(0.018)
Change in	-0.721**	-0.333	-0.707**	-2.407***
Share African-American	(0.320)	(0.321)	(0.314)	(0.685)
Change in	-1.243***	-1.379***	-1.344***	-0.789***
Share Hispanic	(0.255)	(0.261)	(0.262)	(0.259)
	Panel B: Regressions using ACS data			
Change in	0.023***	0.009	0.015***	0.020**
Log Density	(0.006)	(0.006)	(0.006)	(0.010)
Change in	-2.646***	-1.742***	-0.666**	0.017
Share College Education	(0.363)	(0.368)	(0.310)	(0.477)
Change in	0.011***	0.006**	0.006**	0.008
Log Median Income	(0.003)	(0.003)	(0.003)	(0.005)
Change in	4.309	7.312	7.447	14.06
Monthly Rent	(5.091)	(5.587)	(5.581)	(8.723)
Change in	0.025***	0.002	0.001	0.001
Log Median Housing Price	(0.005)	(0.005)	(0.006)	(0.004)
Change in	-0.513**	-0.670***	-0.691***	-0.830
Poverty Rate	(0.219)	(0.227)	(0.227)	(0.533)
Change in	-1.064***	-0.170	-0.135	-0.556
Share Vacant	(0.174)	(0.129)	(0.129)	(0.401)

$***p < 0.01, **p < 0.05, *p < 0.1$ *CD – City Dummies, DV 00 – 2000 Value of Dependent Variable.*

The independent variable is the mean Streetscore change (UCC) of a census tract between 2007 and 2014. Standard errors corrected for spatial correlation are in parentheses. Regressions are estimated with a constant that is not reported.

Table 6: *Is there evidence of Invasion?*

	Coefficient				
Independent Variable	Without CD	With CD	With CD + SS 07 + LD, SCE 00	With CD + SS 07 + LD, SCE 00 (Excl. NY)	With CD + SS 07 + LD, SCE, DTC 00 + Adjacent SS 07
Distance to CBD	-0.013	-0.070***	-0.057***	-0.081***	
	(0.013)	(0.010)	(0.011)	(0.016)	
Adjacent Score	0.142***	0.090***	0.094***	0.090***	
	(0.010)	(0.013)	(0.017)	(0.031)	
Adjacent Log Density	0.376***	0.251***	0.202***	0.322***	0.066
	(0.028)	(0.046)	(0.046)	(0.098)	(0.046)
Adjacent Log Median Housing Price	1.088***	0.590***	0.095	-0.083	0.073
	(0.099)	(0.154)	(0.121)	(0.232)	(0.113)
Adjacent Share College Education	0.011***	0.010***	0.006***	0.007**	0.005***
	(0.002)	(0.001)	(0.002)	(0.003)	(0.002)

$***p < 0.01, **p < 0.05, *p < 0.1$

CD – City Dummies, SS 07 – Mean Streetscore of Census Tract 2007, LD 00 – Log Density 2000, SCE 00 – Share College Education 2000, DTC – Distance to CBD.

All independent variables are from the 2000 U.S. Census. The dependent variable is the mean Streetscore change (UCC) of a census tract between 2007 and 2014. Standard errors corrected for spatial correlation are in parentheses. Regressions are estimated with a constant that is not reported. $***p < 0.01, **p < 0.05, *p < 0.1$

Table 7: *Is there evidence of filtering?*

Independent Variables	SS 07 With CD	SS 14 With CD	SS Change With CD	SS Change With CD
	Coefficient			
1990-2000	0.043***	0.045***	0.002	0.004
	(0.014)	(0.015)	(0.004)	(0.003)
1980-1989	0.046***	0.056***	0.001***	0.006*
	(0.014)	(0.016)	(0.004)	(0.003)
1970-1979	0.020**	0.021*	0.001	0.001
	(0.009)	(0.011)	(0.003)	(0.002)
1960-1969	0.021*	0.024**	0.004	0.002
	(0.011)	(0.012)	(0.003)	(0.003)
1940-1949	-0.003	-0.013	-0.010***	-0.008**
	(0.012)	(0.014)	(0.003)	(0.003)
Before 1940	0.045***	0.051***	0.006***	0.003
	(0.006)	(0.007)	(0.002)	(0.002)
SS 07				-0.002
				(0.014)
Log Density				0.116***
				(0.034)
Share College Education				0.007***
				(0.001)

$***p < 0.01, **p < 0.05, *p < 0.1$ *CD – City Dummies, SS 07 – Mean Streetscore of Census Tract, 2007.* All independent variables are from the 2000 U.S. Census. Standard errors corrected for spatial correlation are in parentheses. Regressions are estimated with a constant that is not reported.