

NBER WORKING PAPER SERIES

COMPUTER VISION AND REAL ESTATE:
DO LOOKS MATTER AND DO INCENTIVES DETERMINE LOOKS

Edward L. Glaeser
Michael Scott Kincaid
Nikhil Naik

Working Paper 25174
<http://www.nber.org/papers/w25174>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2018

We thank Abhimanyu Dubey for his excellent research assistance and acknowledge helpful comments from Jesse Shapiro and Jann Spiess. E.L.G. and N.N acknowledge support from the Star Family Challenge for Promising Scientific Research. 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/w25174.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Edward L. Glaeser, Michael Scott Kincaid, and Nikhil Naik. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Computer Vision and Real Estate: Do Looks Matter and Do Incentives Determine Looks
Edward L. Glaeser, Michael Scott Kincaid, and Nikhil Naik
NBER Working Paper No. 25174
October 2018
JEL No. C40,G21,R30,Z11

ABSTRACT

How much does the appearance of a house, or its neighbors, impact its price? Do events that impact the incentives facing homeowners, like foreclosure, impact the maintenance and appearance of a home? Using computer vision techniques, we find that a one standard deviation improvement in the appearance of a home in Boston is associated with a .16 log point increase in the home's value, or about \$68,000 at the sample mean. The additional predictive power created by images is small relative to location and basic home variables, but external images do outperform variables collected by in-person home assessors. A home's value increases by .4 log points, when its neighbor's visually predicted value increases by one log point, and more visible neighbors have a larger price impact than less visible neighbors. Homes that went through foreclosure during the 2008-09 financial crisis experienced a .04 log point decline in their appearance-related value, relative to comparable homes, suggesting that foreclosures reduced the incentives to maintain the housing stock. We do not find more depreciation of appearance in rental properties, or more upgrading of appearance by owners before resale.

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

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

Michael Scott Kincaid
Department of Economics
Harvard University
mkincaid@fas.harvard.edu

1 Introduction

The value of appearance has been a fundamental axiom in fields like architecture and urban design, but it is now a testable hypothesis. Computer vision techniques enable us to measure a home’s external appearance and to determine how much appearance is valued by the market. A home’s appearance can also be treated as an outcome variable, and we test whether different incentives, created by foreclosure, homeownership and the desire to resell, increase investment in external appearance.

In Section II of this paper, we review the basic methods of visual recognition techniques and discuss our Boston-based data. We use data on the majority of home sales in Boston from 2003 to 2015, and the full set of assessors’ data. We have current images from Google Street View and images from 2007 for essentially all of these homes. We have also been able to use more modest numbers of internal images from Zillow. Our core algorithm follows standard computer vision techniques, where machine learning is used to assess the relative importance of the vast vector of pixel-based data on the appearance of an image.

Location and year of sale explain over 30% of the variation in home sale prices in our data. To focus on the added value of in-person visits and exterior images, we first regress prices on those variables and work with the residual values.¹ On their own, the visual images can explain 11.7% of the price variation of location-residualized prices. By contrast, the assessor-based measure of exterior condition produces an r-squared of .032.

Yet neither measure of external appeal comes close to the explanatory power of standard housing attributes, such as square footage, year build, and the number of bathrooms, which collectively explain 31% of the variance. If we control for these core characteristics, the extra r-squared added by either in-person assessments or images is modest. When the external images are added to these standard variables, the r-squared rises to .328. Adding the assessor’s external measure increases r-squared to .313. Collectively, the images and assessors’ measure increase r-squared to .332. We find roughly similar results when we train the model on one set of the data and examine the explanatory power in a second set of the data.

While these results suggest that neither in-person assessments, nor Street View images, contribute massive explanatory power, they do not imply that appearance doesn’t matter. Figure 1 shows a bin-scatter plot of average price (residualized for location and year of sale) on the visually-predicted price. A one standard

¹[Hough and Kratz \(1983\)](#) began this literature by showing that in downtown Chicago new commercial buildings that won architectural awards commanded higher rents in Chicago; architectural quality did not benefit older buildings. [Vandell and Lane \(1989\)](#) find that commercial buildings that were rated highly by architects in Boston and Cambridge commanded higher rents. [Asabere et al. \(1989\)](#) find a price premium for certain architectural styles in Newburyport, Massachusetts.

deviation rise in the visually predicted price measure is associated with a .16 log point increase in the home's value, or about \$68,000 at the sample mean. This strong association suggests that our visual approach does capture a significant component of housing value, even though visual images do not explain a large share of the unexplained variance of housing prices.

We also tested the natural hypothesis that external appearance is a luxury good. Appearance can explain more of the residualized prices in the richest half of neighborhoods (r-squared of .23) than in the poorest half of neighborhoods (r-squared of .03), when we don't control for other home characteristics. While value is far more visible from the street in rich neighborhoods, the added r-squared from images, on top of standard housing characteristics, is still small in rich neighborhoods.

While an architectural focus emphasizes the link between a building's appearance and its value, urban design emphasizes the link between a neighborhood's appearance and its value. To test the possibility that aesthetic quality spills over to neighboring homes in section IV, we regress one home's price on neighboring home prices, using the images as an instrument for the assessed value of neighboring homes. The estimated coefficient is driven by the correlation of a home's price and the images of its neighbors. The two-stage least squares procedure essentially scales that coefficient so it can be interpreted as standard elasticity. We estimate that the elasticity of home value with respect to the visual component of neighboring prices is .39, which suggests that a home's value increases by 4% if a neighbor's visual predicted price increases by 10%.

The appearance of neighboring homes may be capturing other omitted attributes of the local area, including average income or safety. To focus on appearance itself, we differentiate between nearby homes that are on the same street, and homes that are geographically proximate but not on the same street, and are therefore presumably less visible. The impact of the value of visually conspicuous neighbors is higher than less visible neighbors, when we use our visual instruments. With ordinary least squares, the correlation is stronger with less visible neighbors. Visible neighbors also have a greater impact when we instrument using tax assessor's appearance related variables. These findings suggest that at least part of the correlation between price and neighbor's appearance is working through appearance itself.

In Section V, we use the visually predicted price as a measure of home investment and maintenance. To perform this exercise, we use homes sold between 2006 and 2008 to perform our price prediction exercise. We do this both using only the image, and using the image and the home characteristics. We then use this prediction model to calculate home values, based entirely on image from 2006 onward. We interpret the changes in this measure as capturing the changes in visual quality of the home.

We first show that when after a home receives a permit for remodeling, that home's visually-predicted price rises on average by \$16,000. We see this as a test of whether our measure actually does capture investment in the physical property. We then use our measure of visual upgrading to test three broader hypotheses about incentives and investment in housing. The first hypothesis is that homeowners who are experiencing foreclosure don't maintain their homes. The second hypothesis is that owner-occupiers invest more than renters in their appearance. The third hypothesis is that owners invest before the resale.

The hypothesis that foreclosure leads to depreciation in the housing stock was much debated during the wave of foreclosures that swept over America from 2007 to 2009. The fact that foreclosed homes often sell at a discount is well-known, but that may reflect a depressed market or the desire of foreclosing entities to sell quickly ([Campbell et al., 2011](#)). Images allow us to measure whether the housing actually decays in ways that are visible from the street.

We start with a sample of 1,256 homes that have experienced foreclosure between 2007 and 2010. We then match these homes, using standard nearest neighbor propensity score methods, with a set of homes that do not fit this description. We calculate the visually predicted prices for both sets of homes, using an algorithm trained only on 2007 data. To create that algorithm, we exclude foreclosed homes to avoid any effect that foreclosure might have on price. We calculate the visually predicted price, using this 2007 model, for the homes based on images prior to 2008 and after 2011. We find that the visually predicted price of foreclosed homes fell by nearly \$9000 relative to similar homes that were not foreclosed. We find comparable results using either propensity score methods or standard regressions. If that coefficient generalized to the national level, this coefficient implies that the almost 8 million foreclosures between 2007 and 2009 could have reduced the value of housing by over \$70 billion.

We then test two added hypotheses about homeownership and home maintenance. [Shilling et al. \(1991\)](#) find that price appreciation of owned houses is lower when the home is rented out, perhaps because renters take worse care of their homes. We test this hypothesis by comparing the changes in the visually predicted price for owned and rented homes in Boston. We do not find any significant difference in the change in the visually predicted price of the two types of homes. Perhaps, external investments in homes are easily contractible and so less prone to the moral hazard that may bedevil less obvious actions that maintain home quality.

We also test the hypothesis that resold homes receive added investments in home quality. Such investment would compromise standard repeat sales indices, because if resold homes receive extra investment than

the presumably would appreciate by more than other homes, which did not receive such investment.² We start with a set of homes that are sold between 2013 and 2016. We then use propensity score matching to create a comparable set of homes that did not sell between those years. The mean difference in the change in visually predicted price between the two sets is not statistically different from zero. We also find no significant difference between the resold and not resold homes in regression with a wide range of controls.

Section VI concludes. Visual measurement of housing quality is now possible. Images do not explain a large fraction of the unexplained variance in housing prices, but they do better than variables gathered by in-person assessors. Moreover, the visually-predicted price does seem to provide an index of the value placed on the appearance of a home, which can then be used as an outcome variable. We find evidence for spillovers in appearance across homes and that foreclosed homes depreciate in appearance. We did not find evidence that owned homes improve more in appearance or that appearance improves before resale.

2 Description of Visual Measurement Algorithm and Data

We now discuss our two primary sources of data, Boston Housing Records and Google Street View. We then discuss the technical details of a computer vision-based algorithm for estimating housing prices.

2.1 Boston Housing Records

The city of Boston keeps records on all houses and condominium units within the city, which includes structural characteristics, dates of sale and sale prices. We have addresses which enable us to geolocate each one of the housing units. While we reproduce all of our results with condominium data, we will be focusing on single family homes. This choice is not innocuous. The image data is far more predictive of price with single family homes than with condominiums.

While we have the complete set of homes in Boston, only a fraction of these homes sell during any given year and in some cases the sale price is not recorded by the assessor. Moreover, we must restrict ourselves to sale years in which Google Street View is available. Consequently, we are restricted to examining the 13,221 homes that sold in Boston between 2006 and 2015. Some of these homes sold more than once.

Table 1 provides the means of the institution variables for this sample. The average log sales price in 2016

²While the standard Case and Shiller indices do exclude homes that have engaged in formal remodeling, it is possible that owners do engage in less extensive care before selling which could still bias the repeat sales indices upwards.

dollars is 12.9 (which is equivalent to \$395,000) and the standard deviation of log price is .57. There is a great deal of variation in value of housing in Boston, and the distribution is skewed to the right.

The next four lines show latitude, longitude, year built and year remodeled. The oldest home in the data was built in 1752; the newest in 2014. Remodeling refers to the receipt of a permit to engage in significant remodeling work. The average log of living area is 7.6, which corresponds to about 2,000 square feet. The average log land area is 8.26, which corresponds to 3,800 square feet. There is considerable variation in both interior and exterior space across Boston homes.

The next rows show the number of floors, the number of rooms, the number bedrooms, bathrooms, half bathrooms, kitchens and fireplaces. The range of sizes is quite large, with many smallish houses and a few true mansions. One home in the sample appears to have 21 full bathrooms. We also have variables are produced by the assessor including an aggregation that include categorical variables for outside and inside finishes, building styles, kitchen style, as well as variables that classify the overall condition of interior, exterior and the overall house. These are the variables that reflect the on the ground work of the assessor. There is one index that captures the exterior condition of the house and a second index that captures the interior condition.

2.2 Visual Data from Google Street View and Zillow

To capture the exterior appearance of properties, we obtain building façade images from Google Street View. For each house in our sample, we acquire the time-series of Street View images captured between 2007 and 2016. Since Google updates their imagery every few years for a large city like Boston, we have between two to seven images for all properties in our sample. Each Street View image is associated with a month-year timestamp. In sale price regression experiments, we use the Street View image whose timestamp is closest to the date of sale. In experiments with incentive-related events, we use images captured just before and just after the event has occurred.

To capture the interior appearance of properties, we obtain images from property listings on Zillow, a popular real estate website. We used Zillow’s Application Program Interface (API) to query property records using address and date of sale and obtained listing information including property characteristics and all owner-contributed images. However, property images are typically removed from Zillow after the sale, and hence we were able to obtain images for only 2754 homes in our dataset. Since images posted on Zillow contain interior property views, exterior property views, and floor plans, we developed image recognition algorithms

to remove exterior property images (since we already have these from Street View) and floor plan images. We further used image recognition to identify whether the interior images were of a living room, a bedroom, or a kitchen, since each property record contained multiple views of each of these room types. We then represented each property with three images, one each from a living room, bedroom, and kitchen. While this may not capture the interior appearance of a home in its entirety, it provides a somewhat uniform representation for interior appearance across our small sample. Figure 4 (first row) shows an example from our dataset.

2.3 Computer Vision

Each image in our dataset is a 400x300x3 pixel-resolution color image. To reduce this image to a feature vector, we used the Resnet-101 (He et al., 2016) convolutional neural network (CNN), trained on the Imagenet object classification dataset. The Resnet-101 CNN performs a series of nonlinear operations on the 400x300x3 image tensor to reduce it to a 1024-dimensional feature vector at the penultimate layer. To reduce the dimensionality further, we use principal component analysis and obtain a 100-dimensional feature vector from the top 100 principal components. In sum, each image is represented with a 100-dimensional feature vector, which roughly encodes the shapes, textures, and colors present in the image. These image features are obtained from a neural network trained to process image pixels to predict whether the image contains objects such as cats, dogs, and cars (and not directly trained to predict sale prices from image pixels).

3 Measuring Household Appearance

Our measure of the quality of housing appearance is utterly pecuniary: an attribute of a home’s appearance is deemed attractive if and only if it is associated with higher market prices. We will differentiate between unconditional visually-predicted price and conditional visually-predicted price. Unconditional visually-predicted price is based solely on whether some aspect of the home’s image predicts its sale price unconditionally, with no other items in the regression. Visually-predicted price is based on whether an aspect of the home’s image predicts its sales price conditionally, holding constant everything else about the home.

If we are interested in measuring the pure impact of appearance on either a house’s value or the value of neighboring homes, then surely the conditional visually predicted price is the right measure. If a house is

more valuable because it is bigger, we should not attribute that value to appearance. If we are interested in using visual images to measure investment in the home, the situation is less clear, since we may well be interested in investments that are both related to appearance and also to more general investments. Consequently, we will generally prefer the conditional visually-predicted price when looking at the impact of appearance on price and the unconditional visually-predicted price when looking at the impact of incentives on investment.

In this section, we will generate our measures of the visually-predicted price of the home. We interpret this VPP as an index of the aesthetic appearance of the home, which can then be used to answer questions, including whether attractive homes generate spillovers to neighbors and whether foreclosures generate depreciation in appearance. Following [Mullainathan and Spiess \(2017\)](#), we begin by first orthogonalizing our price data on location and year of sale. Table 2 provides that specification which includes controls for latitude longitude, neighborhood fixed effects and sale year fixed effects. The overall r-squared from that specification is .31, and we view the residuals from that specification as our object of interest. In Table 3, we compare three different approaches to predicting housing prices. In this case, we work with a subset of 16,417 observations, which include those homes for which we have a full set of data and a sale, and a valid sale price. The first regression shows a standard housing price hedonic regression, which controls for the administrative data on rooms, square footage, and year of built, as well as dummies for building style, structure class and exterior finish. The r-squared from this specification is .31, which is roughly comparable with a total r-squared from attributes and location of .5, which is fairly standard for hedonic housing regressions.

The next regression shows the r-squared from using exterior images alone. The .117 r-squared is considerably less than the r-squared from standard features, but it is still far from negligible. The third regression shows the fit of the index of exterior conditions based on the assessor's visit. The assessor's index is strongly correlated with price, but it has a far lower r-squared: .032. The visual images are far less predictive of price than the home attributes, but far more predictive than the assessor index. The fourth regression combines the assessor variable to the basic characteristics and the r-squared rises to .313. The combination of images and basic characteristics produces an r-squared of .328. All the variables together increase the r-squared to .332.

While the images only add modestly to the goodness of fit in standard regression setting, they are strongly associated with price. Figure 1 uses the visually-predicted price, for each home, generating using the model from Table 3, Regression 2, and assuming the mean value for the other characteristics. The figure shows a bin-scatter relationship between visually predicted price and actual price.

A better assessment of the explanatory power generated by the visuals can be done by training prediction models and testing them out of sample. In Table 4, we first train three different models on one half of the data and then test the model on the other half. If the testing sample were example the same as the training sample, then the coefficient on the predicted price should equal one. In the first regression, where we train the model exclusively on the visual data, the estimated coefficient is .826, suggesting that there was overfitting in the training sample. The r-squared for this visual-only index is .092, meaning that almost 10% of the variation in price (residualized first with respect to neighborhood) can be explained by the image model.

The second regression repeats this exercise with the basic variables only. The coefficient is .916 and the r-squared rises to .284. Once again, the basic variables explain far more than the visuals. The third regression trains first on both the basic and the visual variables. The algorithm, naturally, allows for a wide range of interactions between the two sets of variables. The overall r-squared climbs to .285. Once again, explanatory power comes from the core attributes of the house, not from its appearance.

In Table 5, we repeat the exercise in Table 4, but splitting the sample by the income of the block-group. This essentially divides richer neighborhoods from poorer neighborhoods within Boston. We then train two different models. The first uses homes in the poorer neighborhoods. The second starts with homes in the richer neighborhood.

The first panel trains only on the basic variables. The r-squared for the poorer neighborhoods when the model trains on poorer neighborhoods is .16. The r-squared for the richer neighborhoods when the model trains on richer neighborhoods is .46. The ability to predict prices in the richer neighborhoods is much higher than the ability to predict prices in the poorer neighborhoods. When the poor neighborhood trained algorithm is used to predict prices in richer neighborhoods, it has an r-squared of .39, suggesting that the relationship between basic variables and price is similar in rich and poor neighborhoods, and that the relationship between price and these variables is higher in rich neighborhoods.

The second panel includes visuals. In this case, predictive power is low in all but one case. The algorithm trained in poor neighborhoods does a poor job of predicting prices in poor and rich neighborhoods. The algorithm trained in rich neighborhoods does a poor job predicting prices in poor neighborhoods. But the algorithm trained in rich neighborhoods has an r-squared of .23 in rich neighborhoods. Curb appeal is much more strongly related to price in high income neighborhoods.

The final panel shows the impact of adding visuals to the training algorithm. The r-squared for the algorithm trained and tested on poorer neighborhoods rises from .1568 to .1651, a quite modest improvement in

predictive power. The increase in r -squared for the algorithm trained and tested on the rich neighborhood is even smaller. The increase is from .4614 to .4622. The off-diagonal r -squareds actually fall when the visuals are included, meaning that the visual attributes that predict price in the rich sample add noise in the poor sample and vice versa.

Without other data, the visuals are far more predictive of price in rich areas than in poor areas, which suggesting perhaps that visible appeal may be a luxury good. The basic variables have much more predictive power in rich areas than in poor areas, which means that adding the visuals in the poor area generates slightly more r -squared. The visual data adds relatively little r -squared in either data sample.

Table 6 examines interior images taken from Zillow. In this case, we downloaded images from Zillow's website, where we found 2754 of the structures that sold in our data. In this case, we again use residualized prices, and control for the basic characteristics described in Table 3. The r -squared with those controls is .38. The second regression shows the impact of the assessor's index of interior condition, which causes the r -squared to rise to .412. Regression (3) shows that using Zillow images produces an r -squared of .434.

Regression (4) includes both the interior and exterior assessors' indices. The r -squared is .412. Regression (5) includes all of the interior and exterior images, and the total r -squared increases to .429. Regression (6) includes the full set of variables and the r -squared climbs to .449. On their own, the images appear to outperform the assessors' variables, but neither adds that much to the explanatory power and the best fit occurs when we include them all.

As a purely prediction exercise, the images produce a mixed record. They are certainly quite significant predictors of price, but they add only modestly to the r -squared. This may be less true in other settings, where exteriors tell more about the overall quality of the house. But at least in Boston, appearance is far from everything. We now turn from prediction to using the visually predicted price as a measure of the visual quality in other settings. We first ask whether there are spillovers from the appearance of neighboring homes.

4 Do Aesthetics Spill Over?

Cities engage in many policies that manage their aesthetics. Zoning rules restrict the form of buildings. Often neighborhood approval is needed before a building can be permitted. In historic districts, commissions often micro-manage any alterations to the visible part of a structure. These interventions suggest that

urbanites care about the appearance of buildings that surround them. In this section, we test that hypothesis by looking at the price impact of attractive neighboring buildings.

In the previous tables, we restricted our focus just to images of the house itself. We now turn to the impact of neighboring homes. We will test the hypothesis that houses sell for more when they have more attractive neighbors. We cannot, however, rule out the possibility that the better looking neighboring houses are correlated with omitted attributes of either the primary home or the neighborhood, but we will examine whether the appearance of more visible neighbors matters more.

Table 7 begins with our basic neighbor spillover regressions. We perform two specifications. First, we regress the value of a home on the assessed value of the homes of its four closest neighbors, where we instrument for the assessed value of the neighbors using their visual attributes. We use assessed price rather than realized price because it is available for a wider range of homes in a wider range of years.

In this procedure, we are not using appearance as an instrument for neighbor's assessed value because this will test some generic spillover model. Instead, we use this two stage least squares procedure because it implicitly normalizes the impact of visuals in a way that provides an easy interpretation. Intuitively, the two stage least squares estimator, with no other covariates equals $\text{Cov}(\text{Price}, \text{Neighbors' Appearance}) / \text{Cov}(\text{Neighbors' Appearance}, \text{Neighbors' Price})$, which also equals $\text{Cov}(\text{Price}, \text{Neighbors' Appearance}) / \text{Var}(\text{Neighbors' Appearance})$ divided by $\text{Cov}(\text{Neighbors' Appearance}, \text{Neighbors' Price}) / \text{Var}(\text{Neighbor's Appearance})$. The two stage least squares estimate is roughly equivalent to the ratio of the coefficient of price regressed on neighbors' appearance divided by the coefficient of neighbors' price on neighbors' appearance. Consequently, the coefficient provides us with the elasticity of a homes price with respect to the visually predicted price of its neighbors.

The estimated coefficient for the four closest neighbors is .387 which is estimated with a standard error of .012. This coefficient has the interpretation that as the predicted price of a home's four closest neighbors increases by 10%, the sales price of the home itself increases by about 4%. This effect seems significant both economically and statistically.

The largest challenge for interpreting this finding is that this correlation might reflect omitted variables about the neighborhood. A neighborhood with richer people might have better aesthetics. The value of the neighborhood could easily come from its wealth, rather than the quality of the paintwork on the buildings.

To address these concerns, we split nearest neighbors based on street address. We then separately regress on the prices for the four closest neighbors that are on the same street with the four closest neighbors that are

not on the same street. We assume that both sets of neighbors are equally informative about the quality of neighborhood, but the homes on the same street are more visually important to the home.³ Consequently, if aesthetics themselves matter then we expect to see that the coefficient for the same street neighbors is higher than the coefficient for the different street neighbors.

We estimate a coefficient of .399 for the own street neighbors and .274 for the different street neighbors. There is a larger correlation between price and prices from neighbors on different streets, but a larger correlation between price and visually predicted price for neighbors on the same street. This suggests that our approach may be actually capturing something about the neighborhood aesthetics.

In regressions (4), (5) and (6) of the table, we perform the same key two stage least squares procedure using visual variables from the assessor's office rather than variables from Google Street View. These variables include a grade for exterior condition (1–5), a categorical variable indicating building style (e.g., Victorian), a categorical variable indicating structure class (e.g., single-story), and a categorical variable indicating exterior finish (e.g., brick). All four variables encode information about neighboring houses that is visible from the street.

The results are somewhat stronger. The 2SLS coefficient for the 4 nearest neighbors, regardless of street, is .56. The two stage least squares coefficient is about .61 for same street neighbors and .31 for different street neighbors. Once again, there seems to be a spillover from aesthetic appearance.

These results are potentially useful in thinking about a variety of zoning and land use controls. If our coefficients reflect the true treatment effect of near neighbors' appearance, rather than omitted variables, then they suggest that there is a neighborhood level externality. One home's upkeep impacts its neighbors. This finding explains why neighborhood associations often attempt to control local aesthetics, and predicts that events that lead to depreciation of a single home, perhaps including foreclosures, may impact the larger neighborhood.

Table 8 examines the link between assessed value and the visually predicted price of the neighbors homes. In this case, the measured connection is even stronger. Certainly, assessors also seem to believe that having attractive neighbors increases value. The large coefficient in this regression, relative to Table 7, suggests that they may even be overestimating the importance of neighbors' appearance.

Table 9 shows the basic correlates of observable upgrading. In this table, we regress the change in the visually predicted price between 2009 and 2014 on initial values. The first coefficient shows a striking degree

³We are grateful to Jesse Shapiro for suggesting this test.

of mean reversion in the variable, which reflects a combination of true mean reversion (more investment in places that were initially less nice) and spurious mean reversion generated by measurement error. Density and education are both positively correlated with visual upgrading, which is a pattern that we found in our previous work on measuring changes in the visual score of the neighborhood (Naik et al., 2017).

Holding density and education constant, median income is negatively associated with visual upgrading and percent white is positively associated with change in the visually predicted price. In all cases, these variables are quite statistically significant and have economically meaningful magnitudes. For example, a 10 percentage point increase in the share of the population with a college degree is associated with a .01 increase in visually predicted price or a one percentage point increase in observable value.

5 Remodeling, Homeownership, Foreclosure, and Changes in Housing Appearance

In this section, we turn from the consequences to the causes of changes in housing images. We focus on the incentives of homeowners' to keep up appearances. We hypothesize that owners have stronger incentives to maintain looks than renters. Owners may have particularly strong incentives to improve curb appeal before they sell the home. During a foreclosure, owners may have particularly weak incentives to maintain appears and perhaps even to destroy value out of spite.

We examine the impact of resale, homeownership and foreclosure on housing appearance, as measured by the visually predicted price. We also examine whether appearance improves after remodeling which is a basic test of our methodology. In all cases, we will compare samples of homes that are "treated" with control samples that are generated using nearest neighbor propensity score methods. We will supplement these with standard regression results. The identification concerns are somewhat different in the three cases, so we discuss them sequentially.

Our outcome variable is the change in the visual aspect of price between 2009 and 2016. We estimate visually predicted price using a training sample of 1,000 homes with sales prior to 2009. Our first method includes no controls, beyond geography. Our goal is to capture any visual change in the home, not just those changes that are orthogonal to listed characteristics. Consequently, we worked with the broader measure of visual quality.

Figure 2 illustrates the relationship between change in visually predicted price and the sale price of prop-

erties in 2009. A clear u-shaped pattern is visible. There is the most significant upgrading with expensive properties that are presumably bought by richer Bostonians who want their homes to be even nicer. But there is also significant upgrading in places with the lowest sale prices in 2009, which suggests that some low end homes are being bought as “fixer-uppers”.

Figure 3 shows the relationship between appearance in 2009 and change in appearance between 2009 and 2016. Here the pattern shows strong mean reversion until we reach the most attractive homes. Homes with the least attractive appearance in 2009 seem to have received the most attention, which is compatible with a filtering model in which homes first run down and then are fixed up again. Mean reversion does not appear for the homes with the highest indices; those homes experienced visual index upgrading that was higher than expected.

These two figures illustrate how a price based visual index can be used to measure the physical change in the neighborhood, and to establish basic facts about the patterns of home investment and maintenance. We next test whether this measure actually matches remodeling permits, which is really just evidence we are capturing the physical attributes of the home (Table 10, experiments 1-2). We then turn to whether foreclosures are associated with physical decline (Table 10, experiments 3-6), and whether resold homes have physical upgrades, which is important for the interpretation of repeat sales indices (Table 11, experiments 1-3). We end by measuring whether renting is associated with external depression (Table 11, experiments 4-6).

5.1 Housing Appearance and Remodeling

In our sample, 1,025 out of our 16,147 homes received permits for significant remodeling work between 2010 and 2013 from the City of Boston. Our first test is whether these homes also experienced a meaningful increase in their appearance as measured by the visually predicted price, based on difference in the appearance of images captured in 2009 and 2014. Remodeling could include small innovations, like upgrading the clapboard front, or large remodeling, like adding a floor. Figure 4 (second row) shows two examples of remodeling from our dataset.

If the visually-predicted price actually reflects the value of a structure, then presumably it should increase when an owner invests in that structure. As remodeling is not a single treatment, we can, at best, capture the average treatment impact of remodeling across Boston homes during this time period. Moreover, since owners who remodel may have taken a number of smaller actions to maintain appearance, there results are probably best seen as reflecting the impact of having an owner who remodels, not the impact of any single

remodeling action.

The main identification concern is that the remodeling decision is correlated with exogenous factors changing the home's appearance. For example, if remodeling is more likely among homes that are experiencing faster external decay, then the treatment effect from a naïve difference-in-difference estimator should be biased downwards, since the remodeled homes would have been particularly bad without the remodeling. If remodeling is associated with benign positive shocks to appearance, then the treatment effect may be biased upwards. To address these concern, use propensity score matching to create a comparison group.

Many standard identification concerns in real estate are not concerns here. For example, we would expect to see more remodeling in communities that are experiencing exogenous increases in the level of demand. But there is little reason to suspect that increases in demand will improve the appearance of a home unless the owner action takes some action, so this should not bias our result.

We estimate effects for all homes, and for single family units alone. For the sample of all homes, we have 1025 treated homes and 5576 homes in our control sample. We have 424 treated single-family homes, and a control sample of 2073 single family homes.

For the full sample that requested remodeling permits, the mean increase in visually predicted price is .017 log points. There was no change in the control sample. For the single family homes, the remodeled homes received an average visually predicted price increase of .046 log points. The non-treated homes experienced an average increase in visually predicted price of .006 log points.

These mean differences, however, do not control for initial differences in the two sets of homes. To control for such differences, we will match each treated home with five homes in the control sample using two propensity score based methods. Our first method generates a propensity score based exclusively on the visual index in 2007. Our second method generates a propensity score based on the visual index, home characteristics and neighborhood. In both cases, we use nearest neighbor matching for five neighbors.

The first row of this experiment finds that the remodeled homes experienced an increase in visually predicted price of .03 log points, which is about \$15,971 at the sample mean. When we match based on initial images, location and other attributes, the treated effect increases to .041, which is \$21,949 at the sample mean, where the sample refers to the treatment group. The fact that the estimated treatment effect increased with the stronger set of controls suggests that there is negative selection into remodeling. Homes were more likely to be remodeled when they were similar to homes that were depreciating more rapidly, which seems quite plausible. People remodel more when the alternative—not remodeling—gets worse. Somewhat surprisingly,

when we duplicate these results in a regression on the third line, controlling for initial visual index and other characteristics, we do not find a significant remodeling effect on visually predicted price, which we suspect indicates the greater value of the more flexible propensity score approach in this context.

The next rows look at single family attached homes only. The effect of remodeling is positive, but it is not statistically significant when we look at homes that are matched based only on visuals. The effect becomes much larger and significant, when we look at homes that are matched based on visuals and other characteristics. The estimated coefficient is .049, which corresponds to a value increase over \$26,000 at our sample mean. Once again, the regression based result is not significant, but it is closer in magnitude to the other two estimates.

These results are not inherently surprising. We should certainly expect remodeling to increase the visual appearance of the home, and this serves mainly to test our methodology. Since there is great heterogeneity in the nature of remodeling, it is almost meaningless to talk about a single treatment effect from remodeling a house.

5.2 Foreclosure and Changes in Appearance

[Campbell et al. \(2011\)](#) examine forced sales caused by either a bankruptcy or a death. They find that forced sales are associated with a significant decline in home value, which they interpret as reflecting poor maintenance by older owners or neighborhood vandalism. Their work has been interpreted as suggesting that foreclosures could also depreciate the housing stock, again because owners choose not to maintain or protect their homes. Our next exercise in this section is to examine the impact of foreclosures on visual appearance.

We examine the sample of 1256 homes that experienced foreclosure between 2007 and 2010 in Boston, using images from 2007 and 2011. Figure 4 (third row) shows two examples of foreclosures from our dataset. These homes are unlikely to be random, and their physical structures may have been more likely to decay even without the impact of foreclosure. To address this problem, we yet again use propensity score matching to create a sample that is similar in visual appearance, neighborhood and other characteristics to the foreclosed homes.

As discussed above, we measure visually predicted price based on an algorithm using only sales from the pre-period for 1000 homes that are not included in the subsequent estimates. We use the specification where there are no other home level controls other than location. We match our 1256 foreclosure homes with 3601

other homes. The mean change in visually predicted price for the foreclosure homes is $-.018$ log points. The mean change for the control homes is $.003$ log points.

We use propensity score matching to match each member of the foreclosure sample with five members of the control sample. We first do this using only the visual data about the homes in 2007. We then match based on a wider range of home characteristics, including structural characteristics and location.

When we match only on initial image, the estimated impact of foreclosures is $-.023$ log points, which is equivalent to a \$5962 decline in visual value at the mean for the foreclosed sample. When we match based on other home characteristics as well, the estimated impact of foreclosure rises to $-.03$ log points, which represents a \$7750 loss in visual value after foreclosures. The ordinary least squares based estimate is somewhat smaller ($-.014$), but it is still statistically distinct from zero. These effects are statistically significant and economically meaningful, but they are not huge. Often estimates suggest that foreclosures destroy more than 20% of the value of the home. Our estimates of the loss in externally visible value are much smaller, possibly because much of the depreciation is interior and possibly because the actual depreciation was less than prior estimates. Moreover, this loss should not be confused with a welfare loss since, presumably, maintaining the non-foreclosed homes required some effort and cost as well.

The next sample looks at the subset of foreclosed homes that did not also undergo remodeling, as measured by receiving a permit. We see this exercise as purging those homes that may have been explicitly remodeled because of the foreclosure process and focusing on the change in visual value associated with normal wear and tear. We have 890 such homes. Our control sample falls to 2788 homes, and we match using propensity scores, again first based only on initial visuals and then based on visuals and other characteristics.

When we match only on prior visual image, the estimated decline in value is $-.024$ log points, which roughly equivalent to the $-.023$ log points estimated for the complete sample. When we match using the broader range of characteristics, we estimate an impact of $-.035$ log points, which is equivalent to a \$9103 loss in value. The ordinary least squares coefficient is $-.019$. When we look exclusively at the normal wear-and-tear on homes that did not undergo remodeling, foreclosure seems to have slightly larger effects.

In the next panel, we look only at single family homes, which may be more representative of the larger U.S. housing stock. Our sample of foreclosed homes drops to 363. Our control sample falls to 1987. The foreclosed homes in this sample experienced a decline in visually predicted price of $.021$ log points. The non-foreclosed homes experience an increase in visually predicted value of $.017$ log points.

Our first propensity score match, which uses initial images along, estimates a treatment effect of $-.032$ for

foreclosure, which is slightly less than the difference in means. When we match based on a larger set of characteristics, we estimate a larger treatment effect of $-.04$, which represents a loss of \$8924 at the mean for this subsample. The ordinary least squares coefficient is $-.039$, which is also significant and close to the other estimates.

In the last panel, we look only at single family homes that have not been remodeled. The estimated treatment effects increase in magnitude to $-.038$, based on the visual only propensity score, $-.047$, when using the augmented propensity, and $-.037$ when using ordinary least squares. These estimates suggest that single family homes may lose almost 5% of their visual value when they go through foreclosure and do not experience remodeling.

We cannot speak to the generalizability of results from the City of Boston. There are reasons why these results might either be higher or lower than the country as a whole. The relatively harsh New England climate could also make the depreciation of more significant in Boston than elsewhere in the country. Conversely, in Boston, structure represents a smaller share of home value, relative to location, and saw a 5% decrease in structure value in Boston might easily reflect a larger percentage change somewhere else.

5.3 Resale and Changes in Appearance

Standard indices of home prices, such as the Case-Shiller Index and the Federal Housing Finance Agency Index, rely on changes in resale prices: the growth in sales price between a home sold today and some years in the past. Such indices are potentially biased upwards if resold homes receive extra investment before they go on the market. A desire to sell could provide strong incentives to invest in attractive appearance. In that case, the price increase includes both the rise in the market price and the extra impact of the investment. Conceivably, there could also be a reverse bias if homes that had depreciated particularly quickly were more likely to be resold.

Standard indices understand this threat and typically exclude homes that have explicitly undergone remodeling. Yet it is possible that resold homes have been improved in more subtle ways that did not involve an explicit remodeling job. We will test this hypothesis by comparing the change in visually predicted prices for homes that were resold during the 2010-2013 period with homes that were not resold, using images from 2009 and 2014.

Our treatment sample includes 572 homes that were resold during this time period. Our control sample begins with 5024 homes. The resold homes on average saw a decline in $.016$ in visually predicted price.

The homes that were not sold experience an increase of .001 log points.

We first match using visual images only, and estimate an impact of resale of -.018 log points, which is not statistically distinct from zero. We then match based on a wider range of characteristics and estimate an effect of -.021 log points, which is on the edges of statistically different from zero. The ordinary least squares estimate of -.016 is also on the borderlines of significance. These results may mean that homes with other problems were more likely to be resold or that owners remodel when they expect to stay in their homes, but it certainly doesn't suggest an upward bias in the home index.

We then split our sample of resold homes into those homes that experienced remodeling and those homes that did not. We are particularly interested in the homes that did not experience remodeling, since those are the ones that would actually be in a repeat sales index.

The second panel looks at the unremodeled homes. In this case, the homes that were resold experienced more decline in the visually predicted price. However, the effect is statistically insignificant and economically small. The third panel looks only at single family homes that were resold. In this case, the sample for Boston was small. Consequently, the standard errors of our estimates are unfortunately large. The point estimates are, however, quite close to zero, except for the ordinary least squares estimates which is strong, negative and somewhat anomalous.

We interpret these results to suggest that there does not seem to be upward bias in the repeat sales index because of externally visible investments by resellers. However, we cannot reject the possibility that there is internal investment that is not observable in our external images. Moreover, our standard errors are sufficiently large that we cannot rule out a modest amount of upward bias in the indices.

5.4 Homeownership and Changes in Appearance

Our final "treatment" is owner-occupancy. Homeowners might have stronger incentives to invest in the appearance of their homes than renters. [Shilling et al. \(1991\)](#) found that housing values appear to depreciate more quickly when homes are rented out, which suggests that renters may perform less maintenance. We test their hypotheses by examining whether visually predicted price depreciates more quickly for rented homes, using occupancy-status data of homes between 2010-2013 and their images from 2009 and 2014.

Yet the prediction that renting should reduce investment in maintenance is most plausible when maintenance involves unobservable effort that is most easily supplied by the resident. In that case, the landlord cannot

easily write a contract compensating the tenant for care, but external appearance does not need to be supplied by the resident. Many important tasks involving external appearance can be supplied by professional contractors, or potentially by the landlord himself. As such, it is quite possible that renting may reduce internal quality of home, but not reduce appearance.

There are two different challenges to identification within this exercise. First, it is possible that owned properties received more appearance related maintenance for reasons other than ownership. Second, it is possible that ownership is correlated with exogenous shocks to appearance, such as homes that are experiencing faster decay or experiencing more abuse from the external forces. For example, owned homes may be better built in unobserved ways and consequently may be decaying less quickly.

We have 4307 owner-occupied homes in our sample. We match this with a control sample of 5137 units that are occupied by renters. Across the entire sample, the mean growth in the visually predicted price for owner-occupied homes is .006. The mean growth for rented homes is .013.

When we match based on visual data only propensity scores, we estimate a difference of less than .001 for the two sample and this is measured with sufficient precision that we are able to reject any significant difference between the two matched samples. When we match based on the broader set of covariates, we estimate a treatment effect of .006 log points, with a standard error of .008. Taken literally, this estimate does suggest a positive impact of homeowner, but it is less than one-tenth of 1% per year.

Our next panel examines the set of homes that were remodeled to test the hypothesis that owners invest more in remodeling efforts. In this case, the rented units received larger increases in visually predicted price overall. When we match based on visual data, the effect disappears and becomes effectively zero. When we match based on more variables, there is a slight positive effect, but it is statistically insignificant.

When we use a regression, the coefficient on ownership becomes strongly negative. One possible interpretation of this result is that landlords cares more about curb appeal to attract prospective tenants, while owners care more about their own use. Consequently, owner-occupiers remodel in ways that are less visible than landlords.

In line with the previous literature, we now turn to single family detached homes. Multi-unit dwellings are fundamentally different, both because they can house both an owner and a renter and because they may face the slightly different difficulty of coordinating the actions of multiple owners within a single structure. Yet, when we look at a sample of 856 owner occupied single family homes and compare them with a control sample of 2716 rented single family homes, we find almost no effect at all. The same results hold when we

look only at single family homes that have been remodeled, except that again in the ordinary least squares results, ownership again seems to have negative effects.

The ownership results provide little support for the view that ownership increases maintenance, at least on the outside. Perhaps landlords are prone to invest in exteriors to raise rents. We cannot see anything about investment in interiors.

6 Conclusion

This paper, examined the ability of images to predict housing prices in Boston. On their own, the images can explain almost 12% of the variation in price, controlling for location, but the price-related information in images is mostly captured by the standard variables for home size and characteristics. To us, this suggests at least in Boston, you can't tell the value of a home by its exterior, but our point estimates do also imply that better appearance significantly increases price.

We also use interior images to add more explanatory power. Together interior and exterior images raised the r-squared of the hedonic regression from 38 to 43%. It seems reasonable that the interior condition is less likely to be captured by standard housing characteristics.

We tested whether the appearance of neighbor's impacts value. The basic results show a strong spillover from the visually predicted price of a neighbor's home. As the appearance-related value of the eight nearest neighbors' homes increases by .1 log points, a home's value increases by .07 log points. To test the possibility that these spillovers reflect omitted neighborhood characteristics, we compare the impact of neighboring homes that are on the same street with neighboring homes that are on different streets. We found that the appearance of neighboring homes on the same street had a larger impact on price, suggesting that at least some of this correlation reflects aesthetic spillovers.

One use for image-based data is to test how image impacts price. A second use is that images can capture changes in the condition of a home. Using images from 2007 and 2014, and a pricing algorithm based on 2007-era sales, we calculated the change in the visually predicted price for homes in Boston. Reassuringly, we found that price increases were larger for homes that had applied to receive permits for remodeling.

We also found that greater depreciation in homes that went through foreclosure from 2007 to 2009. Our preferred specification is that foreclosed homes lost about 4% of their visually observable value from 2007 to 2014 relative to non-foreclosed homes. These estimated effects are far lower than some estimates of the

social cost of foreclosure, but still significant both economically and statistically.

Our other results are more negative. We don't find that the rented homes depreciate more quickly than owned homes based on their visual images. We don't find that homes that are about resold experience more improvement in observable imagery. The latter result suggests that repeat sales indices are not biased because resellers disproportionately upgrade their homes.

There is considerable potential for future work bringing computer vision techniques into real estate economics. We have attempted to show two avenues for future research: the impact of appearance on price and the use of images to test hypotheses about home maintenance. There are surely others, and we believe our results badly need duplication in other markets and during different time periods.

References

- Asabere, P. K., G. Hachey, and S. Grubaugh (1989). Architecture, historic zoning, and the value of homes. *The Journal of Real Estate Finance and Economics* 2(3), 181–195.
- Campbell, J. Y., S. Giglio, and P. Pathak (2011). Forced sales and house prices. *American Economic Review* 101(5), 2108–31.
- He, K., X. Zhang, S. Ren, and J. Sun (2016). Deep residual learning for image recognition. pp. 770–778.
- Hough, D. E. and C. G. Kratz (1983). Can “good” architecture meet the market test? *Journal of Urban Economics* 14(1), 40–54.
- Mullainathan, S. and J. Spiess (2017). Machine learning: an applied econometric approach. *Journal of Economic Perspectives* 31(2), 87–106.
- Naik, N., S. D. Kominers, R. Raskar, E. L. Glaeser, and C. A. Hidalgo (2017). Computer vision uncovers predictors of physical urban change. *Proceedings of the National Academy of Sciences* 114(29), 7571–7576.
- Shilling, J. D., C. Sirmans, and J. F. Dombrow (1991). Measuring depreciation in single-family rental and owner-occupied housing. *Journal of Housing Economics* 1(4), 368–383.
- Vandell, K. D. and J. S. Lane (1989). The economics of architecture and urban design: some preliminary findings. *Real Estate Economics* 17(2), 235–260.

Table 1: Summary Statistics ($N = 16417$)

	Mean	Std. Dev.	Min	Max
Log Sale Price	12.885	0.571	10.768	17.086
Location				
Latitude	42.307	0.037	42.232	42.392
Longitude	-71.097	0.042	-71.174	-70.996
Basic Features				
Year Built	1917.307	30.228	1752	2014
Year Remodeled	2001.22	12.257	1925	2014
Year Built (Normalized)	0.631	0.115	0	1
Year Remodeled (Normalized)	0.423	0.491	0	1
Log Living Area	7.668	0.433	6.086	9.603
Log Total Land	8.259	0.618	5.973	11.375
Num. Floors	2.178	0.624	1	4
Num. Bedrooms	4.622	2.025	1	16
Num. Full Bathrooms	2.055	0.942	1	21
Num. Half Bathrooms	0.387	0.575	0	10
Num. Kitchens	1.708	0.806	0	4
Num. Fireplaces	0.434	0.857	0	12
Assessor Exterior Features				
Exterior Condition	2.438	0.814	1	5
Assessor Interior Features				
Interior Condition	2.715	0.972	1	5

*Notes: In addition to the variables listed above, the following categorical variables are included: The **Location** feature group includes a Neighborhood dummy. The **Basic Features** include Air-conditioning Type, Heating Type, Owner-occupied Flag, Building Style, Exterior Finish, and Structure Class. The **Assessor Interior Features** include Interior Finish, View, Kitchen Style, and Bathroom Style.*

Table 2: *Sale Price and Location*

	(1) Log Sale Price
Latitude	11.194*** (0.162)
Longitude	-3.965*** (0.130)
Neighborhood Fixed Effects	Yes
Sale Year Fixed Effects	Yes
Observations	16417
Adjusted R^2	0.310

Standard errors in parentheses

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

Table 3: Residualized Sale Price and Property Characteristics

	(1) Basic	(2) Street View	(3) Assessor Exterior	(4) Basic + Assessor Exterior	(5) Basic + Street View	(6) Basic + Assessor Exterior + Street View
Year Built (Normalized)	0.103*** (0.039)			0.108*** (0.039)	0.054 (0.040)	0.057 (0.040)
Year Remodelled (Normalized)	0.016** (0.007)			0.010 (0.007)	0.021*** (0.007)	0.015** (0.007)
Log Living Area	0.236*** (0.016)			0.228*** (0.016)	0.239*** (0.016)	0.233*** (0.016)
Log Total Land	0.108*** (0.008)			0.112*** (0.008)	0.087*** (0.009)	0.090*** (0.009)
Num. Floors	0.133*** (0.012)			0.130*** (0.012)	0.119*** (0.012)	0.116*** (0.012)
Num. Bedrooms	-0.017*** (0.003)			-0.016*** (0.003)	-0.013*** (0.003)	-0.011*** (0.003)
Num. Full Bathrooms	0.077*** (0.006)			0.073*** (0.006)	0.076*** (0.006)	0.073*** (0.006)
Num. Half Bathrooms	0.063*** (0.006)			0.061*** (0.006)	0.061*** (0.006)	0.059*** (0.006)
Num. Kitchens	-0.014 (0.009)			-0.008 (0.009)	-0.008 (0.009)	-0.001 (0.009)
Owner Occupied Flag	0.085*** (0.007)			0.085*** (0.007)	0.082*** (0.007)	0.082*** (0.007)
Num. Fireplaces	0.098*** (0.004)			0.095*** (0.004)	0.089*** (0.004)	0.086*** (0.004)
Exterior Condition			0.104*** (0.004)	0.040*** (0.004)		0.041*** (0.004)
Building Style	Yes	No	No	Yes	Yes	Yes
Construction Type	Yes	No	No	Yes	Yes	Yes
Exterior Finish Type	Yes	No	No	Yes	Yes	Yes
Exterior Image Features	No	Yes	No	No	Yes	Yes
Observations	16417	16417	16417	16417	16417	16417
Adjusted R^2	0.310	0.117	0.032	0.313	0.328	0.332

Notes: All price variables are in log dollars, residualized on location. Standard errors are in parantheses.

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

Table 4: *Residualized Sale Price and Visual Index*

	(1)	(2)	(3)
	Vis. Index	Basic Index	(Basic + Vis.) Index
Vis. Index	0.826*** (0.027)		
Basic Index		0.916*** (0.016)	
(Basic + Vis.) Index			0.874*** (0.015)
Observations	8617	8617	8617
Adjusted R^2	0.092	0.284	0.285

Notes: This table reports the regression coefficient and R^2 values for out-of-sample residualized sale price prediction for different feature groups. Standard errors are in parantheses.

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

Table 5: *Goodness-of-fit across income groups*

	Basic		Street View		Basic + Street View	
Train	< Median ($N = 4247$)	> Median ($N = 4247$)	< Median ($N = 4247$)	> Median ($N = 4247$)	< Median ($N = 4247$)	> Median ($N = 4247$)
Test						
< Median ($N = 4139$)	0.1568	0.1247	0.0335	0.0152	0.1651	0.1216
> Median ($N = 4203$)	0.3886	0.4614	0.0297	0.2267	0.3324	0.4622

Notes: This table reports R^2 values for out-of-sample residualized sale price prediction, where the samples are split by median household income at the block-group level as reported by the 2010 US Census.

Table 6: Residualized Sale Price and Property Characteristics (including Interior Imagery)

	(1) Basic	(2) Assessor Interior + Basic	(3) Zillow+ Basic	(4) Assessor Interior + Exterior+ Basic	(5) Street View + Zillow + Basic	(6) All
Year Built (Normalized)	0.134 (0.089)	0.142 (0.087)	0.173* (0.093)	0.146* (0.087)	0.190* (0.098)	0.186* (0.096)
Year Remodelled (Normalized)	0.024 (0.016)	-0.003 (0.016)	0.023 (0.016)	-0.003 (0.016)	0.030* (0.016)	0.007 (0.017)
Log Living Area	0.324*** (0.037)	0.281*** (0.036)	0.342*** (0.039)	0.281*** (0.036)	0.332*** (0.040)	0.299*** (0.040)
Log Total Land	0.097*** (0.018)	0.101*** (0.018)	0.095*** (0.019)	0.102*** (0.018)	0.084*** (0.023)	0.094*** (0.023)
Num. Floors	0.141*** (0.026)	0.109*** (0.025)	0.120*** (0.027)	0.109*** (0.025)	0.115*** (0.027)	0.092*** (0.027)
Num. Bedrooms	-0.007 (0.006)	-0.002 (0.006)	-0.006 (0.006)	-0.001 (0.006)	-0.005 (0.006)	-0.001 (0.006)
Num. Full Bathrooms	0.030*** (0.011)	0.021* (0.011)	0.027** (0.011)	0.020* (0.011)	0.024** (0.011)	0.014 (0.011)
Num. Half Bathrooms	0.046*** (0.013)	0.038*** (0.013)	0.041*** (0.014)	0.038*** (0.013)	0.037*** (0.014)	0.032** (0.014)
Num. Kitchens	-0.028 (0.019)	0.005 (0.019)	-0.011 (0.020)	0.006 (0.019)	-0.011 (0.021)	0.015 (0.021)
Owner Occupied Flag	0.047*** (0.015)	0.048*** (0.015)	0.028* (0.016)	0.047*** (0.015)	0.024 (0.017)	0.028* (0.016)
Num. Fireplaces	0.105*** (0.009)	0.067*** (0.010)	0.102*** (0.010)	0.066*** (0.010)	0.091*** (0.010)	0.061*** (0.010)
Interior Condition		0.045*** (0.009)		0.042*** (0.009)		0.034*** (0.010)
Exterior Condition				0.018* (0.009)		0.012 (0.010)
Building Style	Yes	Yes	Yes	Yes	Yes	Yes
Exterior Finish Type	Yes	Yes	Yes	Yes	Yes	Yes
Construction Type	Yes	Yes	Yes	Yes	Yes	Yes
Interior Finish Rating	Yes	Yes	Yes	Yes	Yes	Yes
Kitchen Style	Yes	Yes	No	Yes	No	Yes
Bathroom Style	Yes	Yes	No	Yes	No	Yes
View Rating	Yes	Yes	No	Yes	No	Yes
Interior Image Features	No	No	Yes	No	Yes	Yes
Exterior Image Features	No	No	Yes	No	Yes	Yes
Observations	2754	2754	2754	2754	2754	2754
Adjusted R^2	0.382	0.412	0.434	0.412	0.429	0.449

Notes: All price variables are in log dollars, residualized on location. Standard errors are in parantheses.

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

Table 7: Aesthetic Spillovers from Neighbors using Assessed Value as Endogenous Variable

	(1) 4 NN (Images)	(2) 4 NN (Images) Same Street	(3) 4 NN (Images) Different Street	(4) 4 NN (Assessor)	(5) 4 NN (Assessor) Same Street	(6) 4 NN (Assessor) Different Street
First-stage F-stat	533.192	46.595	158.103	77.725	82.25	34.44
2SLS	0.387*** (0.012)	0.399*** (0.047)	0.274*** (0.071)	0.560*** (0.062)	0.607*** (0.033)	0.309 (0.228)
OLS	0.725*** (0.016)	0.734*** (0.071)	0.755*** (0.137)	0.821*** (0.037)	0.807*** (0.104)	0.948** (0.239)
First Stage; Basic Features	Yes	Yes	Yes	Yes	Yes	Yes
First Stage; Neighbor Exterior Image Features	Yes	Yes	Yes	No	No	No
First Stage; Neighbor Exterior Condition Rating	No	No	No	Yes	Yes	Yes
Observations	1546	765	765	1546	765	765

Notes: This table reports the 2SLS estimates of the log sale price, residualized on location. The endogenous variable is average assessed value of 4 neighbors from 2014, instrumented with average features of neighbor's street view image features in columns (1–3) and average features of neighbor's exterior condition, building style, structure class, and exterior finish in columns (4–6). Additional controls are included as shown. Standard errors are clustered by neighborhood.

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

Table 8: *Assessed Value on Average Visual Index of Four Neighbors*

	Log Assessed Value
Avg. vis. index, 4 nearest neighbors	.956*** (.057)
Latitude/Longitude	Yes
Basic Features	Yes
Neighborhood FE	Yes
Own Image Features	Yes
Observations	1551
Adjusted R^2	0.7866

Notes: This table reports the OLS regression of assessed value on average visual index of four neighbors (controlling for location, basic features, and own image features). Additional controls are included as shown.

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

Table 9: 2009–2016 Visual Upgrading on Census Tract Demographics

	Visual Index Diff 2009–2016
Visual index 2009	-.536*** (.013)
Log density	.013*** (.003)
Log median household income	-.020*** (.006)
Percent white	.001*** (.001)
Percent college-educated	.001*** (.000)
Latitude/Longitude	Yes
Basic Features	Yes
Neighborhood FE	Yes
Observations	9121
Adjusted R^2	0.196

Notes: This table reports the OLS regression of the difference in visual index from 2009–2016 (visual upgrading) on demographic variables from US Census (controlling for location, basic features, and initial appearance). Additional controls are included as shown.

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

Table 10: Effect of Economic Events on Visually-predicted Property Price (Remodeling and Foreclosures)

Matching Model	Treatment (#Samples)	Control (#Samples)	Treatment (After-Before)	Control (After-Before)	Coefficient	Standard Error	Z-score
(1) Effect of Remodeling on Visually-predicted Price							
Vis. Index (DID)	1025	5576	0.017	-0.000	0.030***	0.011	2.75
Vis. Index + Basic Features (DID)					0.041***	0.001	3.47
Vis. Index + Basic Features (OLS)					0.003	0.009	0.34
(2) Effect of Remodeling on Visually-Predicted Price (Single Family Homes)							
Vis. Index (DID)	424	2073	0.046	0.006	0.024	0.018	1.35
Vis. Index + Basic Features (DID)					0.049***	0.019	2.60
Vis. Index + Basic Features (OLS)					0.025	0.016	1.53
(3) Effect of Foreclosures on Visually-predicted Price							
Vis. Index (DID)	1256	3601	-0.018	0.003	-0.023***	0.008	-2.81
Vis. Index + Basic Features (DID)					-0.030**	0.009	-3.16
Vis. Index + Basic Features (OLS)					-0.014**	0.007	-1.99
(4) Effect of Foreclosures on Visually-predicted Price (Unremodeled Homes)							
Vis. Index (DID)	890	2788	-0.022	0.004	-0.024***	0.009	-2.60
Vis. Index + Basic Features (DID)					-0.035***	0.011	-3.18
Vis. Index + Basic Features (OLS)					-0.019**	0.008	-2.35
(5) Effect of Foreclosures on Visually-predicted Price (Single Family Homes)							
Vis. Index (DID)	363	1987	-0.021	0.015	-0.032**	0.014	-2.22
Vis. Index + Basic Features (DID)					-0.040***	0.016	-2.46
Vis. Index + Basic Features (OLS)					-0.039***	0.012	-3.28
(6) Effect of Foreclosures on Visually-predicted Price (Single Family Unremodeled Homes)							
Vis. Index (DID)	281	1537	-0.021	0.017	-0.038***	0.015	-2.49
Vis. Index + Basic Features (DID)					-0.047***	0.019	-2.36
Vis. Index + Basic Features (OLS)					-0.037***	0.013	-2.85

Notes: All price variables are in log dollars, residualized on location. In the Vis. Index matching model, a propensity score is constructed on the basis of a home's visually-predicted log price (based on Street View features). In the Vis. Index + Basic Features model, a neighborhood dummy, log living area, year built (normalized), and owner occupied flag is added to the set of covariates used for matching. The table reports the Abadie-Imbens standard errors.

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

Table 11: Effect of Economic Events on Visually-predicted Property Price (Resales and Owner-occupancy)

Matching Model	Treatment (#Samples)	Control (#Samples)	Treatment (After-Before)	Control (After-Before)	Coefficient	Standard Error	Z-score
(1) Effect of Resale on Visually-predicted Price							
Vis. Index (DID)	572	5024	-0.016	0.001	-0.018	0.011	-1.63
Vis. Index + Basic Features (DID)					-0.021*	0.012	-1.72
Vis. Index + Basic Features (OLS)					-0.016*	0.009	-1.78
(2) Effect of Resale on Visually-predicted Price (Unremodeled Homes)							
Vis. Index (DID)	405	3701	-0.020	-0.003	-0.016	0.010	-1.25
Vis. Index + Basic Features (DID)					-0.009	0.017	-0.70
Vis. Index + Basic Features (OLS)					-0.014	0.011	-1.26
(3) Effect of Resale on Visually-predicted Price (Single Family Homes)							
Vis. Index (DID)	166	1306	-0.008	0.014	0.011	0.021	0.55
Vis. Index + Basic Features (DID)					-0.003	0.023	-0.15
Vis. Index + Basic Features (OLS)					-0.034***	0.012	-2.71
(4) Effect of Owner-occupancy on Visually-predicted Price							
Vis. Index (DID)	4307	5137	0.006	0.013	-0.000	0.000	-0.63
Vis. Index + Basic Features (DID)					0.006	0.008	0.83
Vis. Index + Basic Features (OLS)					-0.012*	0.006	-1.78
(5) Effect of Owner-occupancy on Visually-predicted Price (Remodeled Homes)							
Vis. Index (DID)	941	1111	-0.001	0.035	0.000	0.000	0.63
Vis. Index + Basic Features (DID)					0.008	0.017	0.48
Vis. Index + Basic Features (OLS)					-0.054***	0.017	-3.10
(6) Effect of Owner-occupancy on Visually-predicted Price (Single Family Homes)							
Vis. Index (DID)	855	2716	0.016	0.023	0.000	0.000	0.51
Vis. Index + Basic Features (DID)					0.006	0.007	0.83
Vis. Index + Basic Features (OLS)					-0.018	0.012	-1.56
(7) Effect of Owner-occupancy on Visually-predicted Price (Single Family Remodeled Homes)							
Vis. Index (DID)	177	659	0.026	0.049	0.000	0.000	-1.23
Vis. Index + Basic Features (DID)					-0.011	0.031	-0.39
Vis. Index + Basic Features (OLS)					-0.066**	0.029	-2.26

Notes: All price variables are in log dollars, residualized on location. In the Vis. Index matching model, a propensity score is constructed on the basis of a home's visually-predicted price (based on Street View features). In the Vis. Index + Basic Features model, a neighborhood dummy, log living area, year built (normalized), and owner occupied flag is added to the set of covariates used for matching. In the owner-occupancy experiments (4–7), the owner occupied flag is excluded.

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




Figure 1: Binscatter Plot for log location-residualized sale price versus visually-predicted price (#bins = 50)




Figure 2: Binscatter Plot for growth in visual index between 2009-2016 versus log sale price of properties in 2009 (#bins = 50).




Figure 3: Binscatter Plot for growth in visual index between 2009-2016 versus visual index of properties in 2009 (#bins = 50).

Exterior Image



Interior Images



Living Room



Kitchen



Bedroom

Sale Price = \$347,000, Gross Area = 4,759 sqft, Living Area = 1,010 sqft, #Bedrooms = 3



Before remodeling (2009)



After remodeling (2014)



Before remodeling (2009)



After remodeling (2014)



Before foreclosure (2007)



After foreclosure (2011)



Before foreclosure (2007)



After foreclosure (2011)

Figure 4: *Example images from the dataset*