

The Effects of Federal “Redlining” Maps: A Novel Estimation Strategy

Disa M. Hynsjö* Luca Perdoni†

DECEMBER 16, 2022

[Click Here for The Most Recent Version](#)

Abstract

This paper proposes a new empirical strategy to estimate the causal effects of 1930s federal “redlining” – the mapping and grading of US neighborhoods by the Home Owners’ Loan Corporation (HOLC). Our analysis exploits an exogenous population cutoff: only cities above 40,000 residents were mapped. We employ a difference-in-differences design, comparing areas that received a particular grade with neighborhoods that would have received the same grade if their city had been mapped. The control neighborhoods are defined using a machine learning algorithm trained to draw HOLC-like maps using newly geocoded full-count census records. For the year 1940, we find a substantial reduction in property values and homeownership rates in areas with the lowest grade along with an increase in the share of African American residents. We also find sizable house value reductions in the second-to-lowest grade areas. Such negative effects on property values persisted until the early 1980s. Our results illustrate that institutional practices can coordinate individual discriminatory choices and amplify their consequences.

*PhD Candidate, Department of Economics, Yale University. Disa passed away in Summer 2021. This paper is founded on her vision, brilliance and courage.

†Early Career Researcher, School of Economics, University of Edinburgh. Email: luca.perdoni@ed.ac.uk, website: <http://lucaperdoni.com>. I am indebted to my committee members, Joseph Altonji, John Eric Humphries and Cormac O’Dea for their guidance, support and encouragement. Chelsea Ilarde provided exceptional research assistance. For detailed comments and suggestions, I also thank Julian Aramburu, Jaime Arellano-Bover, Paula Calvo, Alberto Ciancio, Rebecca Diamond, Jose Antonio Espin-Sanchez, Paul Goldsmith-Pinkham, Sun Kyoung Lee, Trevon Logan, Costas Meghir, Sarah Raviola, Jaehee Song, Ebonya Washington, Seth Zimmerman. Conference participants at the NBER Summer Institutes on Urban Economics and the Development of the American Economy, the CESifo Area Conference on Labor Economics and the IEB Workshop on Urban Economics have provided precious feedback. The technical support of Barbara Esty and Miriam Olivares from Yale University Library, as well as Dan Feenberg and Mohan Ramanujan at NBER, has been fundamental. This research has been supported by the Washington Center for Equitable Growth and the Cowles Foundation Program in Labor and Public Economics.

1. Introduction

Discrimination in the housing market can take many forms. Racial covenants, restrictive zoning and barriers to credit access are only a few examples. While developers, real estate agents and bank executives have often been blamed for their exclusionary norms, public institutions took the lead in shaping discriminatory practices in some cases. Between 1935 and 1940, a federal agency – the Home Owners’ Loan Corporation (HOLC) – created *Residential Security Maps* for more than 200 US cities to summarize the financial risk of granting loans in different neighborhoods. Color-coded maps assigned each neighborhood one of four security grades, from A (green) to D (red).¹ Standardized forms attached to the maps (*Area Descriptions*) invariably described the presence of African Americans, Jews, and European immigrants as detrimental to a neighborhood’s grade. In the late 1970s, urban historian K. T. Jackson rediscovered the maps at the National Archives and proposed them as an example of structural discrimination (Jackson, 1980). Since then, the view that the HOLC maps were a source of residential redlining, the systematic denial of mortgages to residents of a community, has steadily gained popularity (Rothstein, 2017; Coates, 2014). Nevertheless, some historians have emphasized the program’s timing and confidentiality to raise doubts about whether a federal data collection program could have strongly influenced the housing market (Hillier, 2003, 2005; Fishback et al., 2021). Quantitative evidence supporting either claim about the maps’ influence is scarce due to the lack of detailed and comprehensive datasets and the non-random assignment of different grades. In this paper, we propose a new strategy to measure the causal effects of a federal government initiative that has been proposed as a symbol of structural discrimination by journalists, activists, academics and presidential candidates.²

This paper estimates the short and long-term causal effects of the Home Owners’ Loan Corporation maps using a new empirical strategy. Our approach exploits an exogenous

1. As an example, a scan of the HOLC map for New Haven, CT is available in Figure 1.

2. As an example, see Rothstein (2017) and Coates (2014). Historical government support of redlining practices has been proposed by President J. Biden and Senator E. Warren as a motivation for their housing plans.

population threshold: only cities above 40,000 residents were mapped. We use a machine learning classification algorithm to draw residential security maps in control cities with a population below the 40,000-resident threshold.³ Using the grades predicted by the classification model, we apply a grouped difference-in-differences design to measure the causal effects of the HOLC maps. Our outcomes of interest are homeownership rates, property values, rent prices and shares of African American residents in 1940. We also analyze the evolution of these outcomes between 1960 and 2010. The estimated causal effects are identified by the differences between neighborhoods in treated cities and neighborhoods that would have received the same security grade if their city had been mapped.

The effects we find in red neighborhoods support the view that HOLC maps reduced access to credit and led to urban disinvestment. In 1940, shortly after the introduction of the maps, we find a sizable reduction in property prices and homeownership rates in D (red) areas, along with an increase in the percentage of African Americans living in those neighborhoods. Property value reductions are also detected in C (yellow) areas.⁴ The negative effects on property prices in yellow and red areas persisted until the early 1980s, shortly after the introduction of legislative measures meant to improve access to residential credit.⁵

The credibility of our approach relies on the performance of the machine learning classification model. To assess its precision, we build a test dataset randomly excluding 25% of neighborhoods from the algorithm’s training procedure. We then compare observed and predicted grades in the test dataset. Our trained random forest algorithm assigns the correct grade to more than 90% of test neighborhoods, and its predicted maps are convincing replicas of those made by HOLC. Even if the model is trained on the complete set of municipalities surveyed by HOLC, including all American metropolises, its performance is robust in cities close to our threshold of interest as well. When we restrict the test dataset to neighborhoods

3. Defining the control group with a machine learning algorithm is an alternative to synthetic control methods. In our case, control units are actual observations grouped by a predictive model replicating an observed classification mechanism.

4. This result is consistent with a finding in Aaronson et al. (2021b) labelled by the authors as “yellow-lining”.

5. The Equal Credit Opportunity Act (1974), the Home Mortgage Disclosure Act (1975) and the Community Reinvestment Act (1977) had the common goal of increasing access to mortgages in neighborhoods previously ignored by financial institutions.

from cities with fewer than 50,000 residents, the overall precision is still above 90%.

Our empirical strategy is possible thanks to a spatial dataset we constructed using the 1910-1940 full-count census records (Ruggles et al., 2020), National Historical Geographic Information System (NHGIS) information (Manson et al., 2021) and CoreLogic property deeds. We clean and impute household addresses for each census decade following best practices from the urban history literature (Logan and Zhang, 2018). Detailed geographic coordinates are assigned to census observations using a state-of-the-art locator. Georeferenced data allow us to match census records with digitized HOLC maps and alternative sources of information to expand our dataset to the years beyond 1940. We include sociodemographic information from the National Historical Geographic Information System (NHGIS) (Manson et al., 2021) along with disaggregated property transaction prices from the CoreLogic deed database to estimate the long-term causal effects of the maps. The resulting dataset covers major US urban areas between 1910 and 2010.

Our paper contributes to a growing literature in economics studying the consequences of HOLC policies. This paper is most closely related to the work of Aaronson, Hartley and Mazumder (2021b), who use a border regression discontinuity design to measure the local effects of lower grades in cities surveyed by HOLC. Unlike Aaronson et al. (2021b), our estimation method compares similar neighborhoods in mapped and unmapped cities. In particular, we avoid spatial discontinuity designs because of endogeneity concerns due to differential pre-trends in socioeconomic variables on different sides of the borders traced by the HOLC. Aaronson et al. (2021b) document the differential pre-trends, due to non-random location of borders and non-random assignment of grades.⁶ They employ propensity scores and a subset of idiosyncratic borders to address endogeneity concerns. Moreover, our empirical approach measures a different type of effect. While the existing literature has focused on the local effects of receiving a lower HOLC grade, we capture the global effects of the four HOLC grades. In our case, the counterfactual is made of similar neighborhoods not mapped by HOLC rather than nearby areas with a higher evaluation. With respect to Fishback, LaVoice, Shertzer and Walsh (2020), who investigates whether HOLC grades were racially-biased, we

6. See also Fishback et al. (2020).

study a different question: the effects of HOLC grades on property prices and demographic characteristics. While there is a growing body of research on HOLC mapping,⁷ we are the first to propose a predictive model replicating the HOLC maps and employing an exogenous population threshold for estimating the effects of different grades.

Focusing on the effects of institutionalizing a set of exclusionary attitudes is an important complement to existing research in the economics of discrimination. Economists have mainly focused on individuals who discriminate based on their taste (Becker et al., 1971), because of imperfect information⁸ or implicit bias (Bertrand et al., 2005; Bertrand and Duflo, 2017). These different mechanisms originate from individual choices and cannot be readily applied to settings where something other than an individual discriminates.⁹ This paper provides an empirical analysis focused on *institutional* discrimination, revealing new evidence of an overlooked source of socioeconomic inequality. The results capture the impact of institutional assessment practices developed by a governmental organization and adopted in the real estate market at large.

Another relevant dimension of the HOLC initiative was its technological content. The agency undertook an unprecedented data collection effort,¹⁰ creating a data analytics tool at the forefront of real estate appraisal techniques of the time. HOLC maps can be interpreted as an innovation in statistical technology that led to increased automation in the processing of mortgage applications. Today, concerns about algorithmic bias (Rambachan et al., 2020; Ludwig and Mullainathan, 2021) and distributional impacts of statistical technology are widespread (Fuster et al., 2021). Our results characterize the effects of a federal initiative

7. All the recent papers on this topic either focus on different questions or employ different empirical approaches and datasets. In addition to Aaronson et al. (2021b) and Fishback et al. (2020), see Fishback et al. (2011), Faber (2020), Aaronson et al. (2021a) and Fishback et al. (2021). There is also a number of contemporaneous working papers on this topic using spatial regression discontinuities designs: see Anders (2019), Appel and Nickerson (2016) and Krimmel (2018).

8. See Fang and Moro (2011) for a review of research on statistical discrimination originated by Phelps (1972) and Arrow (1973).

9. See Small and Pager (2020) and Lang and Kahn-Lang Spitzer (2020) for a comparison of perspectives on discrimination between Sociology and Economics. The sociological literature has focused more on institutional sources of discrimination, if compared with economics.

10. See Michney (2021) for a description, based on HOLC staff correspondence, of how the maps were developed.

that provided a powerful and practical tool to evaluate local housing market conditions.¹¹ They also offer a cautionary tale of how institutional practices can coordinate individual biases and amplify their consequences.

The paper also contributes to the literature investigating the causes of segregation and urban inequality.¹² As outlined in [Boustan \(2013\)](#), residential segregation can result from individual choices by white households,¹³ Black self-segregation,¹⁴ or collective action.¹⁵ Our findings give an example of the last of these causes since the federal agency’s practices had the effect of reinforcing residential exclusion. In terms of methods, we contribute to a relatively recent body of literature using machine learning algorithms to build control groups for causal inference in observational studies.¹⁶

The paper proceeds as follows. Section 2 provides additional details about HOLC activities and the circulation of its maps, while Section 3 contains a description of our novel dataset. We outline our empirical strategy in Section 4, along with results about the performance of our classification algorithm and an array of validity checks. The estimated effects of HOLC maps can be found in Section 5. Section 6 concludes.

2. Historical Background

In the aftermath of the Great Depression, the Roosevelt Administration developed several programs to tackle a mortgage crisis characterized by soaring default rates and falling property

11. In particular, we provide an empirical analysis of the consequences of a collection of federal maps. See [Nagaraj and Stern \(2020\)](#) for a review of recent work about the Economics of maps.

12. See [Glaeser and Vigdor \(2012\)](#), [Cutler et al. \(1999\)](#) and [Logan and Parman \(2017\)](#) for an overview of trends for different urban segregation measures.

13. The mechanism is often referred to as “white flight”. See [Boustan \(2010\)](#) and [Boustan \(2016\)](#).

14. This possible source of segregation does not find strong empirical support. See [Krysan and Farley \(2002\)](#) and [Ihlanfeldt and Scafidi \(2002\)](#).

15. Collective action to induce segregation can take many forms. Some examples are racial covenants ([Jones-Correa, 2000](#); [Sood et al., 2019](#)), urban renewal programs ([Collins and Shester, 2013](#)) and public housing programs ([Chyn, 2018](#); [Tach and Emory, 2017](#)).

16. An example is [Lieberman et al. \(2018\)](#). See [Mullainathan and Spiess \(2017\)](#) for a review of machine learning algorithms within the econometric toolbox.

values. The Home Owners' Loan Corporation (HOLC) was created in 1933 to aid homeowners “in hard straits largely through no fault of their own” (Federal Home Loan Bank Board, 1937). Under the direction of the Federal Home Loan Bank Board (FHLBB), the first task of the HOLC was to refinance mortgages in distress with longer terms, lower interest rates and higher loan-to-value ratios. In particular, the HOLC granted fully amortized loans with 15-year minimum terms at 5% interest rate, financing up to 80% of the property value.¹⁷ The agency concluded its \$3 billion lending effort in 1936 after refinancing over one million loans and holding approximately 10% of US non-farm mortgages (Jackson, 1980).

As a consequence of their lending program, HOLC gained considerable exposure to the housing market. Government officials believed that a healthier lending industry was necessary to safeguard the value of federal real estate investments (Hillier, 2005). In particular, they considered the standardization of appraisal techniques critical to achieving price stability. For this reason, the FHLBB directed HOLC to develop a systematic evaluation process for US neighborhoods, following a growing interest in ecological models across the real estate industry.¹⁸ In 1936, HOLC started the *City Surveys* program, producing maps (*Residential Security Maps*) and standardized forms (*Area Descriptions*) for 239 major U.S. cities. The initiative was completed by 1940.

Field agents drew HOLC maps based on published reports, public records, federal maps and detailed surveys of local financial institutions (Michney, 2021). The availability of these sources varied between cities, and HOLC agents relied on their networks in the real estate community to supply any missing information. The result was meant to be “a composite opinion of competent realtors engaged in residential brokerage, good mortgage lenders and the HOLC appraisal staff.”¹⁹ HOLC agents traced boundaries to divide residential areas into homogeneous neighborhoods. They then assigned a grade on a four-level scale meant to

17. These terms were much more convenient to homeowners than the 5-year interest-only loans, with interest up to 7%, that were prevalent in the market up to that time. See Fishback et al. (2011).

18. See Jackson (1980) and Light (2010) for a discussion about how ecological models, newly-developed by the Chicago School of Sociology, became an influential theory for real estate appraisal.

19. Corwin A. Fergus to T. L. Williamson, October 2, 1936, Roll 431, Home Owners Loan Corporation, microfilm copies of General Administrative Correspondence 1933-36, National Archives II (College Park, MD). As cited in Michney (2021).

summarize the financial security of real estate investment in each zone.²⁰ Areas colored in green (grade A) were the first-tier neighborhoods, while blue neighborhoods (grade B) were deemed still good. The color yellow (grade C) highlighted neighborhoods becoming obsolete or at risk of “infiltration of a lower grade population” (Hillier, 2005). Red neighborhoods (grade D) were considered “hazardous” (Hillier, 2003) for investment. The agency also produced detailed *Areas Descriptions* for each neighborhood. In these forms, they described housing conditions, local amenities and the area’s demographic composition.²¹ The presence of African Americans, Jews and certain European immigrants was inevitably characterized as a “detrimental influence” that “infiltrated” the American social fabric with fatal effects on local housing markets (Jackson, 1980). While the inclusion of racial and ethnic hierarchies in real estate appraisal was pervasive at that time, HOLC practices implemented these notions at an unprecedented scale with the coordinated effort of more than 20,000 employees distributed across more than 200 local offices, and the stamp of federal approval.

HOLC could not have used the results of the *City Surveys* in its lending decisions since the maps were created after the agency completed its refinancing effort. Therefore, the economic impact of the maps depends on how widely these documents circulated among other federal agencies and private financial institutions. The literature offers diverging views on this topic. Hillier (2003) reports that the FHLBB intended to restrict access to “agencies within the FHLB” and “such government agencies having interests allied with those of the Board” while no copies were granted to “private interests”. However, the author concedes that the maps were in strong demand among the public and that local consultants employed by the HOLC had access to these documents. An opposite stance, first proposed by Jackson (1980) and more popular today, argues that HOLC’s findings were widely distributed and quickly became a benchmark for real estate appraisal both in government agencies and the private sector.

Even if we lack definitive evidence about the circulation of *City Surveys*, there is proof that another federal agency – the Federal Housing Administration (FHA) – received multiple

20. A scan of the HOLC map drawn in 1937 for New Haven, CT can be found in Figure 1.

21. As an example, a scan of the HOLC Area Description for New Haven D-4 neighborhood is available in Figure 2.

copies of HOLC maps. The FHA evaluated applications to its mortgage insurance program with manuals that described the presence of “undesirable racial or nationality groups” in a neighborhood as detrimental. Moreover, the FHA employed a collection of maps that categorized neighborhoods on a four-level scale according to their financial security (Aaronson et al., 2021b). Today, a systematic comparison between the HOLC security maps and those of the FHA is impossible, but historical research provides evidence that the two collections were often similar.²² If so, HOLC maps can be considered the best available proxy for FHA standards of neighborhood appraisal.²³ While the HOLC ceased its activities in 1951, the FHA continued its operations in the following decades.²⁴

The explicit inclusion of racial or ethnic criteria in real estate financing became illegal in 1968 with the introduction of the Fair Housing Act. A further series of federal laws enacted in the 1970s addressed concerns about the lasting effects of financial exclusion. In particular, the Equal Credit Opportunity Act (1974), the Home Mortgage Disclosure Act (1975) and the Community Reinvestment Act (1977) were meant to counteract redlining by reinforcing anti-discrimination legislation, introducing mortgage disclosure requirements and supervising credit supply at the local level.

3. Data

We construct a new dataset drawing from three sources: digitized HOLC maps, census data and CoreLogic deeds records. Our classification algorithm is trained on 1930 census data merged with HOLC maps. The short-term effects of the HOLC maps are measured with

22. Nearly all FHA maps are missing. A limited comparison is possible thanks to a reproduction of the FHA map of Chicago, IL (Light, 2010).

23. Fishback et al. (2021) study FHA-backed mortgages in three US cities between 1935 and 1940. The vast majority of loans were granted in areas rated A or B by HOLC maps. However, the authors argue that FHA exclusionary patterns were established before HOLC maps were drawn, and they did not change throughout their study period.

24. There is no historical evidence of different FHA appraisal practices according to the 40,000 resident threshold, or any other population threshold. Moreover, there is no evidence about how the FHA used HOLC maps. We assume homogeneous FHA practices in treatment and control cities, except for the availability of the HOLC maps.

1930 and 1940 census data, while the long-term effects are estimated by combining full-count census data with CoreLogic real property data and tract-level census data for the decades between 1960 and 2010.²⁵

We plan to extend the analysis to 1950 when the full-count census data for that decade becomes available in April 2022. In the following sections, we provide additional details about each data source.

3.1 HOLC Residential Security Maps

We incorporate HOLC grades in our project using the digitized maps provided by the Digital Scholarship Lab at the University of Richmond (Nelson et al., 2021). The files contain maps for 202 cities in 38 states. We convert neighborhood shapes originally traced by HOLC into a regular grid of hexagons. Hexagons are our fundamental spatial unit of observation, and their use simplifies the construction of HOLC maps in control cities.²⁶ The area of one hexagon approximates the typical size of a block in US grid plan cities such as New York City and Chicago.²⁷ We assign a grade to a hexagon if one color occupies at least 75% of its surface.²⁸ This spatial transformation has a negligible impact on the overall distribution of the grades. Appendix Table A1 shows the proportions of each grade according to different spatial definitions, while Figure 3 compares the digitized version of the HOLC map of New Haven, Connecticut, with its hexagon-level counterpart. The percentage reductions for A and B grades are due to the smaller average size of HOLC neighborhoods in these classes. Furthermore, while our hexagons have a fixed area, the HOLC neighborhoods do not, which explains the minor discrepancies in the shares for C and D grades.

25. Validity checks of the empirical strategy employ census data between 1910 and 1930, the pre-treatment decades.

26. More details on why we choose to use a grid of hexagons can be found in Section 4.1.

27. The grid is made of regular hexagons with an area of 0.025 squared kilometers (7.3 acres) and a side of approximately 100 meters (328 feet).

28. The results are robust to modest variations in the 75% threshold. Given the small dimension of each spatial unit, the vast majority of hexagons (81.2%) contains only one grade. 7.5% of hexagons do not meet the 75% threshold and have a missing grade.

3.2 Census Data

3.2.1 Full-Count Census Data

We rely on full-count census records for data between 1910 and 1940. We geocode the heads of household by taking advantage of the addresses available in the proprietary version supplied by *ancestry.com* and IPUMS (Ruggles et al., 2020). Census addresses are cleaned following best practices found in the spatial history literature (Logan and Zhang, 2018), and geographic coordinates are assigned by a state-of-the-art locator (ESRI StreetMap Premium 2019) that combines parcel centroids and street locations.²⁹ Detailed geographic coordinates allow us to construct neighborhood-level averages by combining individual observations with our graded hexagon grid. Population distributions according to the grades can be found in Table 1. Yellow and red areas include 77.8% of the general population in our 1930 sample, but contain 95.6% of African American respondents. Table 2 reports descriptive statistics of the 1930 census according to the HOLC grades. Even before the agency’s intervention, African Americans were concentrated mainly in red neighborhoods. Homeownership rates, property values, and rent prices are all positively correlated with the HOLC scale.

3.2.2 NHGIS Data

Starting in 1950, we must rely on publicly available census data. We obtain tract-level data for homeownership rates, property values, rent prices, and the shares of African Americans between 1950 and 2010 from the National Historical Geographic Information System (NHGIS) at IPUMS (Manson et al., 2021). We focus on census tracts since they are the smallest geographical units identifiable between 1950 and 2010.³⁰ Census tracts became available in smaller cities only in later decades. Hence, this source does not provide full coverage of our sample of interest until 1980.³¹ This is the best available nationwide source of harmonized data for demographic characteristics and homeownership rates in the second half of the

29. The overall proportions of matched addresses for 1910,1920,1930 and 1940 are 60.5%, 65.4%, 76.1% and 73.5% respectively.

30. The median population of a census tract in our dataset is 231, while it is 68 for hexagons. The hexagon area is constant while the one of the census tract is not. A census tract is always bigger than the hexagons we defined in surface terms. The median census tract contains 34 hexagons in surface terms.

31. As mentioned in Section 4, we focus on cities with population between 30,000 and 50,000. Appendix Table A2 reports rates of coverage of NHGIS data for our sample of interest.

twentieth century.

Geographical coordinates allow us to harmonize information from HOLC maps, 1910-1940 full count census data and 1960-2010 NHGIS data. We can use our composite dataset to describe the socioeconomic evolution of US neighborhoods throughout the twentieth century according to the grades assigned by HOLC in the late 1930s. Figure 4 contains trends for our four outcomes of interest between 1910 and 2010, showing that the HOLC ranking in terms of homeownership rates, property values and rent prices was stable during the last century. D (red) neighborhoods were, and still are, the most likely residence for African Americans. The percentage of Black American residents increased in A, B and C areas after World War Two. In particular, C (yellow) neighborhoods reached a 10% share of African Americans in 1970, while B (blue) neighborhoods met the same threshold in 1990. A (green) neighborhood did not attain the same level in 2010 yet.

3.3 CoreLogic Deeds Records

We supplement NHGIS tract-level data with sale records obtained from CoreLogic which contains transaction data collected from county assessors and deed registries, including information about sale prices, dates of sale and the geographic coordinates of the buildings. In our dataset, transactions are binned into 5-year windows according to the sale year and month. As expected, the number of sales recorded in the dataset is much higher in recent years.³² The nature of CoreLogic data is different from data sources we have described so far. They are administrative records of realized sales, while census property values are the results of extensive surveys based on self-reports.

4. A Novel Estimation Strategy

We propose a new strategy to measure the short and long-term effects of the HOLC maps. Our approach does not rely on border discontinuities designs, which have been prevalent

32. Additional details about CoreLogic’s coverage of our cities of interest can be found in Appendix Table A2.

in the literature on the topic.³³ Instead, we exploit an exogenous population threshold: HOLC staff focused on cities with a population of at least 40,000 residents. Accordingly, we define cities with a population between 40,000 and 50,000 as the treated cities, while the municipalities between 30,000 and 40,000 residents are included in the control group. Figure 5 illustrates this threshold and highlights our definitions of treated cities in purple and control cities in orange.³⁴ However, a city-level analysis cannot measure grade-specific effects. Moreover, ignoring the heterogeneous effect of the four different grades would lead to empirical results that would mischaracterize the legacy of this federal intervention. Appendix Table A3 shows the results we obtain when we apply a simple difference-in-difference design to estimate the effect of HOLC mapping imposing homogenous effects for the four different grades. We do not detect any significant effect on homeownership rates or African American percentages, while we find weak evidence of a reduction in property values. To measure the consequences of HOLC maps, we need an empirical approach to estimate the impacts of four different interventions that share the same treatment assignment and timing.

We develop a strategy to compare areas with a given grade in treated cities with neighborhoods that would have received the same grade if their city had been mapped. A machine learning classification algorithm assigns grades to neighborhoods in control cities, replicating HOLC assessment standards. The algorithm is trained to link observed grades from the whole set of HOLC maps to 1930 census observables aggregated at the neighborhood (hexagon) level. Using the predicted grades, we then apply a grouped difference-in-differences design to measure the causal effects of the maps' four different grades. To provide an intuition, let us restrict the analysis to two cities: a treated city, Phoenix, Arizona, and a control city, Raleigh, North Carolina. To estimate the effect of a D (red) grade, we compare the outcomes for observations geocoded in Phoenix D areas with Raleigh observations in neighborhoods

33. The previous literature has mainly focused on the local effect of a lower HOLC grade using spatial regression discontinuity designs. This approach suffers from endogeneity concerns due to the non-random location of borders and non-random assignment of grades. Both [Aaronson et al. \(2021b\)](#) and [Fishback et al. \(2020\)](#) document how locations on opposite sides of HOLC borders showed differential trends in a variety of observables before the introduction of the maps. [Aaronson et al. \(2021b\)](#) employ propensity scores and a subset of idiosyncratic borders to address this issue.

34. See Appendix A.4 for a list of cities.

(hexagons) that the HOLC would have rated D if its agents had surveyed the area.

Building the control group with a machine learning model is a matching technique alternative to synthetic control methods (Abadie and Gardeazabal, 2003; Abadie et al., 2010). Instead of estimating a complete set of weights so that control areas could mirror treated areas, we classify neighborhoods by replicating HOLC standardized evaluations. In the same spirit of synthetic controls, we do not use any post-treatment data when designing the control group classes and the contribution of each observation to the counterfactual is explicit.³⁵ Moreover, the “donor pool” for each class can be easily visualized on a map. Unlike the synthetic control method, in our procedure control group units are never used in the training procedure that determines the counterfactual composition. The resulting control group will have to meet validity checks, such as the parallel trends assumption, that were not targeted during its design.³⁶ Hence, our approach reduces even more the possibility of manipulation in developing a synthetic counterfactual. Harnessing an institutional feature of our research setting, we defined control groups that, while being synthetic, are particularly plausible.

4.1 A Classification Algorithm

The success of our strategy relies on convincingly replicating HOLC evaluations using 1930 census data. In particular, we are interested in recovering a function that can credibly predict y , the HOLC grade, based on X , a set of neighborhood observables.

Since HOLC appraisers traced area boundaries and assigned grades simultaneously, our classification algorithm should replicate both outcomes. We tackle these goals by classifying a regular grid of hexagons into the four different grades.³⁷ This approach imitates HOLC

35. Our approach can be thought of as a special case of the synthetic control method where weights are assigned by a classification model. In particular, for treated observations with grade j we are building a control group assigning weights with only two values, either 0 or 1. Let $\hat{g} \in \{A, B, C, D\}$ be the grade predicted by the classification model, we are assigning the 0 weight to all control observations such that $\hat{g} \neq j$ and a weight equal to 1 if $\hat{g} = j$. The resulting weighted mean will be rescaled by $n_{\hat{g}}$, the number of observations with a predicted class \hat{g} .

36. The machine learning training dataset does not contain information about pre-treatment trends. It only includes 1930 census information.

37. Hexagons, rather than triangles or squares, are well suited for our goals because they are the most circular-shaped polygon that can generate a regular grid. In particular: they reduce sampling bias, capture

methods and tackles the complex task of drawing grade borders in control cities. While the hexagon grid is useful to imitate the original borders, 1930 census data are the best nationwide data source to replicate HOLC surveys. Our dataset provides sufficient detail about the sociodemographic composition of neighborhoods and housing prices. However, we lack information about mortgages, defaults and interest rates that the agency regularly collected surveying local financial institutions.

We implement a random forest algorithm (Breiman, 2001) to classify hexagons into one of four HOLC grades. This machine learning method proves effective in dealing with the class imbalance of our classification problem³⁸ and outperforms other popular classification algorithms.³⁹ In short, the random forest is a nonparametric and nonlinear model based on decision trees. A tree is a hierarchical series of splitting rules for covariates X . In practice, the goal is to find the best binning structure of covariates X , together with the hierarchy of these splits, to predict class y . Since random forests are widely employed in recent economics literature, we will highlight only a few relevant features of the algorithm.⁴⁰

A decision tree provides flexible binning of multiple covariates to maximize the predictive power for the outcome class y . The definition of bins is entirely data-driven, and the process flexibly takes into account interactions between covariates. The resulting bins define a link between covariates X and the predicted class \hat{y} as a nonlinear multivariate function. This approach usually returns a good in-sample fit, but it often suffers from poor out-of-sample predictions due to overfitting. Bootstrapped aggregation (bagging) techniques offer a remedy. The solution is to fit several trees on bootstrapped samples of the data, thus growing a forest. Moreover, each split is determined only by m randomly selected covariates. These steps reduce the correlation between the predictions of each tree, characterizing the forest as “random” and providing reliable out-of-sample predictions. Once the algorithm is trained,

curved patterns more easily, reduce the projection distortion due to earth curvature and provide a better definition of neighbors because of their centroid properties. For more details see Birch et al. (2007).

38. The minority class share (Grade A) is 7.8% while the maximum one (Grade C) is 42.1%. More details can be found in Table A1.

39. More details about the performance of an ordered logit model in this setting can be found in section 4.2.

40. See Fuster et al. (2021) for a more detailed explanation of the random forest algorithm.

the predicted class is the one most voted on by all the trees in the forest. The number of variables to use at each split (m), and the fraction of observations to sample for each tree, are parameters that need tuning.⁴¹

4.2 Classification Model Results

We train the random forest algorithm with a hexagon-level dataset containing all cities mapped by the HOLC with a population below 3,000,000.⁴² The dataset includes 48 different 1930 census variables; some are included at different geographical levels, bringing the total number of training variables to 163.⁴³ The total number of neighborhoods in our dataset is 192,016.

We assess the performance of our classification model on a test set of spatial units (hexagons) randomly drawn from the original dataset.⁴⁴ These observations were excluded from the random forest training procedure and represent an out-of-sample validation of the model performance. Table 3 presents a matrix comparing the observed and predicted grades in the test set (*Confusion Matrix*). The probability of correctly classifying a neighborhood (*Accuracy*) is 91.55%, while the probabilities of correct predictions conditional on observed grades (Class-specific *Sensitivities*) are above 90% for B, C, and D classes. Comparing the predicted grade distribution (*Detection Prevalence*) with the observed class frequencies (*Prevalence*) shows that our model does not alter the overall distribution of HOLC grades. Given that the identification strategy focuses on neighborhoods in cities between 30,000 and 50,000 residents, we are interested in the performance of our model in smaller cities. Appendix Table A4 shows the results we obtain if we restrict our test set to cities with a population below 50,000. Accuracy is still above 90% percent and performance metrics are

41. More details on tuning of our random forest can be found in Appendix section A.3.

42. The results in this section are robust to variations in the population threshold determining which cities are included in the training procedure. The range of variation for this threshold is between 50,000 and 7,000,000 residents.

43. The complete list of variables is available in Appendix A.3.

44. The test dataset represents 25% of the original dataset, while 75% of the observations were used in training the model. The random sampling was stratified according to HOLC grade and city population. The results are robust to changes in the sampling procedure.

similar overall.

It is worth comparing our machine learning procedure to classification models traditionally used in the economics literature. Accuracy levels above 90% are a substantial performance improvement compared to what we would obtain with an ordered logit. Appendix Table A5 shows the performance of an ordered logit estimated on the same dataset used by the random forest. Overall accuracy reaches only 64.35% and the predicted grades severely overestimate the prevalence of C neighborhoods, while underestimating the presence of D and A neighborhoods. It should be noted that a logit-type model is more transparent than a random forest in characterizing the contribution of each variable to determine the probability of a grade. However, the improvement in predictive accuracy offered by the latter is so significant that it compensates for the loss in interpretability.

Prediction accuracy characterizes the model’s precision, but it does not provide any insight into the spatial patterns of our predicted maps. In particular, the challenge is to obtain graded areas with a sufficient degree of compactness to mirror the HOLC maps. A comparison between the original HOLC map of Baltimore, Maryland, and our predicted map can be found in Figure 6. In general, the predicted neighborhoods are not dissimilar from the original neighborhoods in terms of shape.⁴⁵ When the classification model disagrees with HOLC evaluations, it tends to assign a different grade to whole clusters rather than to single hexagons. Examples of this behavior can be found in downtown Baltimore where the predicted grade is C (yellow) versus an original D (red), or in the northwest suburbs of the city, where an area with a B (blue) grade from HOLC is classified as A (green) by the model. The final goal of replicating HOLC grades is to draw “redlining” maps in cities between 30,000 and 40,000 people. Examples of predicted maps for control cities can be found in Figure 7. The model identifies areas for all four grades, returning neighborhood shapes similar to the ones observed in the original maps in larger cities.

The random forest algorithm we employ does not have any spatial constraint that would

45. The surface covered by our hexagon-level maps is slightly smaller than the area originally covered by HOLC. This is because the federal agency mapped even scarcely populated areas, while our strategy focuses on hexagons with at least 20 residents.

guarantee an output visually similar to HOLC maps. The results rely on a training dataset that includes several census observables at different levels of geographical definition. Figure 8 shows how the predicted map for New Haven, Connecticut changes when we train the classification models with datasets including different geographical levels of aggregation. The top left map is the output we obtain when each hexagon only includes information about its area, while the top right map adds city and county-level information to the dataset. These maps, while generally accurate, suffer from spatial noise and the resulting neighborhoods cannot be easily encircled into a compact shape.⁴⁶ The plausibility of the predicted maps increases when we include information about the area surrounding each hexagon. In particular, we construct averages of surrounding census observables using 500-meter and 1,000-meter radii (0.31 miles and 0.62 miles, respectively) for each hexagon. With the addition of these local averages, the classification algorithm returns predicted maps with compact neighborhoods, as shown in panels (c) and (d) of Figure 8.⁴⁷

4.3 Measuring the Effects of HOLC Maps

Our goal is to estimate the grade-specific causal effects of introducing HOLC maps, an innovative information tool for neighborhood appraisal, in the real estate market. In the early stages of the 20th century information revolution, the HOLC maps could act as a coordination device providing practical area evaluations to local financial institutions. To estimate these effects, we classify neighborhoods into four classes according to their predicted grades using our trained random forest. Then, we apply a difference-in-differences design comparing neighborhoods in treated cities, between 40,000 and 50,000 residents, with those in control cities, between 30,000 and 40,000 residents, separately for each grade. Our pre-treatment period is 1930, and 1940 is our first post-treatment period. If the empirical design assumptions are deemed credible, the estimated coefficients will capture the global effects of each grade

46. The overall prediction accuracy of the random forest considering only neighborhood level data is 66.75%, while it rises at 75.45% when including city and county level covariates.

47. A random forest trained on a dataset including neighborhoods and information about their surroundings achieves a 89.58% accuracy. If we add city and county level variables to the former dataset, accuracy increases to 91.55%.

assigned by HOLC.

In the short run, our specification is:

$$Y_{i,h,c,t}^{\hat{g}} = \alpha^{\hat{g}} D_c + \gamma^{\hat{g}} P_t + \beta^{\hat{g}} D_c P_t + \delta^{\hat{g}} X_{i,h,c} + \varepsilon_{i,h,c,t}^{\hat{g}} \quad (1)$$

In the equation, $Y_{i,h,c,t}^{\hat{g}}$ is the outcome for individual i , living in a neighborhood h with grade \hat{g} , in city c , at time t . The term D_c is a treatment dummy and P_t is a post-treatment indicator. $X_{i,h,c}$ includes neighborhood observables and information about the surrounding areas. Equation (1) will be estimated by group according to the grade \hat{g} assigned by the trained random forest. The coefficients of interest are $\{\beta^A, \beta^B, \beta^C, \beta^D\}$. In section 5.1 we provide results for different specifications of equation (1), replacing the treatment term $\alpha^{\hat{g}} D_c$ with city fixed effects or neighborhood (hexagons) ones.

We can extend equation (1) to measure medium and long-term effects. In particular, we estimate the following equation at the neighborhood level:⁴⁸

$$Y_{h,c,t}^{\hat{g}} = \alpha^{\hat{g}} D_c + \Gamma^{\hat{g}} \bar{P}_t + \sum_{t \in T} \beta_t^{\hat{g}} D_c P_t + \delta^{\hat{g}} X_{h,c} + \varepsilon_{h,c,t}^{\hat{g}} \quad (2)$$

where $T = \{1930, 1940, 1960, 1965, 1970, 1975, \dots, 2010\}$. The only new terms compared to equation (1) are $\bar{P}_t = (P_{1940}, P_{1960}, \dots, P_{2010})$ a vector of year dummies for all elements of T (except 1930) and its corresponding vector of coefficients $\Gamma^{\hat{g}}$. In this case, $X_{h,c}$ includes time-invariant geographic controls. Section 5.2 contains the results for different specifications of (2) where we replace the treatment term $\alpha^{\hat{g}} D_c$ with city fixed effects.

The validity of this empirical framework relies on two main assumptions. First, the maps did not affect dependent variables in control cities. Second, outcomes would have evolved in parallel in the absence of the policy. We examine the validity of these assumptions in the following section.

48. We switch to a neighborhood level regression to estimate the long-term effects of HOLC maps. While CoreLogic data allow an individual level analysis with disaggregated deeds, NHGIS data do not. To ensure comparability between the two sources of post-1940, data we will focus on neighborhood level results. The long-term results obtained with the CoreLogic dataset are robust to using an individual-level specification.

A limitation of our analysis is related to external validity. Our effects are estimated for cities with a population between 30,000 and 50,000 and might not be appropriate to describe the effects of the maps in American metropolises.⁴⁹ Another limitation is that our estimates rely on a first-stage classification model. The prediction errors of the random forest might attenuate, or inflate, the difference-in-differences estimates and affect their precision.⁵⁰ Ultimately, the high prediction accuracy of our trained machine learning algorithm and the soundness of the parallel trends assumption reassure us about the overall credibility of this strategy.

Previous research on this topic has focused on estimating the local effects of a lower grade, such as D , compared to a higher one, e.g. C , with border regression discontinuity techniques. Such results characterize within-city local impacts, but it is not immediate to translate them in an aggregate measure of HOLC maps' effects. Instead, our approach returns the treatment effects of the four HOLC grades providing a direct description of the global effects of HOLC maps on US neighborhoods.

4.4 Validation of the Empirical Strategy

As noted in section 4.3, our difference-in-differences framework relies on the no-treatment-spillover assumption and parallel pre-trends. In our context, the no-spillover assumption means that control cities would not have been affected by the HOLC intervention because of their geographic location. Figure 9 shows that control cities, in blue, are scattered throughout the country, and they are not suburbs of treated cities, in red. To strengthen the assumption, we include in our analysis only control cities with a distance of at least 50km (31mi) to the nearest treated municipality.⁵¹ While it is safe to assume that HOLC mapping did not directly affect control cities, it is harder to argue that the assignment of grade g in a certain

49. In Section 5.1 we show results when we expand the treatment group to include cities up to 60,000 inhabitants.

50. Our results are robust to substituting the predicted grades with the observed ones in treated cities. Moreover, we propose additional checks of biases introduced by our classification exercise in section 4.4 to mitigate these concerns.

51. The results are robust to variations in this threshold between 30km (18.6mi) and 70km (43.5mi). The median distance between a control city and the closest mapped municipality is 144.7Km (89.9mi)

area does not affect surrounding neighborhoods.⁵² If we are worried about the spillovers of surrounding graded areas in treated cities, the coefficient β^g will combine the effect of grade g and the correlations with other local grades. In Table 4 we provide the grade composition of neighborhood surroundings according to their own grade and treatment status. D grades are surrounded, on average, by 62.8% red and 32.3% yellow neighborhoods. For all grades, the majority neighborhood share corresponds to the same class: for example, 76.7% of C neighbors belong to grade C. In section 5.1, we show that the results are robust when we include information about the local grade composition as an additional control.

The soundness of the difference-in-differences framework hinges on the similar evolution of socioeconomic characteristics between treatment and control cities prior to the HOLC intervention. We graphically investigate the soundness of the parallel trends assumption in Figure 10 using data between 1910 and 1930, the decades before the HOLC intervention. The trends for African American percentage and homeownership rate evolved in parallel in C (yellow) and D (red) areas. The same is true for grade B, as can be seen in Appendix Figure A1. The assumption appears less valid for A (green) areas, and the results for this class, which represents approximately 3% of the sample, should be interpreted with caution.

This validity check cannot be completed for property values, one of our outcomes of interest, because the census started to record this variable only in 1930. As a partial remedy, we can investigate the trajectory of alternative socioeconomic variables. The bottom panels of Figure 10 compare trends for the imputed income score we built based on 1940 census information.⁵³ Given its definition, this variable can be interpreted as an index of socioeconomic status. The observed trends support our research designs. Additional plots investigating the trends for population density, number of children and percentage of first-generation migrants can be found in Appendix Figure A2. Table 5 contains the results for an analytical check of the

52. This challenging problem is similar to estimating the effects of an exogenous shock when dealing with non-random exposure, as described by [Borusyak and Hull \(2020\)](#). In our setting, even if it is credible to characterize treatment assignment as random, we might think that neighborhood location will lead to non-random exposure to different grades from the surrounding areas.

53. The Census did not record income before 1940. We impute an income score for wage-employed men aged 25-55 between 1910 and 1930 using income measures from 1940. More details can be found in Appendix section A.2.

parallel trends assumption. There are no substantial differences between treated and control cities in the evolution of demographic and economic variables between 1930 and 1920 in B, C, and D areas. These results are confirmed in Table A7 for changes between 1920 and 1910.

The classification algorithm we employ to replicate HOLC grades returns predicted maps that we can compare with the original ones in terms of socioeconomic characteristics. Table 6 compares averages according to observed and predicted grades, showing that the predicted maps do not alter the original socioeconomic composition of C and D neighborhoods. Moreover, we might be worried about the type of bias introduced in the analysis by spatial units receiving a “wrong” grade.⁵⁴ The averages in Appendix Table A6 show that even when the model assigns a neighborhood to the wrong class, we are not introducing significant sources of bias.

Another assumption implicit in our empirical approach is that HOLC practices did not change between different cities. In particular, a predictive model trained with US metropolises might not accurately replicate HOLC grades in smaller cities, the ultimate goal of our prediction exercise. Appendix Figure A3 shows that the accuracy level of our random forest algorithm is robust to different training datasets according to the population of cities included in the training set. When we restrict our attention to predicting grades for neighborhoods in cities below 50,000 residents, overall accuracy is still above 90%. We interpret these results as evidence of the high degree of standardization of HOLC grading procedures, making our predictive model a reliable source for HOLC evaluations in the smaller control cities.

5. The Effects of HOLC Maps

5.1 Short Term Results

We start by estimating equation (1) separately for each grade. The coefficients of interest reported in the tables of this section are $\{\beta^A, \beta^B, \beta^C, \beta^D\}$. The dataset includes individual-level observations from 1930, the pre-intervention period, and 1940, the post-intervention

54. Spatial units with different observed and predicted grades represent approximately 12% of our sample in treated cities.

one. Standard errors are clustered at the city-year level.⁵⁵

Table 7 reports the results for local homeownership rates. The results from a simple difference-in-differences design are reported in column (1). Replacing the treatment indicator with a city fixed effect does not substantially alter the coefficients but considerably increases the estimates' precision. Instead, we do not gain additional precision if we substitute a city-level fixed effect with a neighborhood (hexagon) fixed effect as in column (3). The last column reports our preferred specification where we combine a city fixed effect with neighborhood level sociodemographic controls. We find a 2.4 percentage points decrease in the percentage of homeowners in D (red) zones in 1940, shortly after the introduction of the maps. A weaker effect can be detected in C (yellow) areas, while we find no effects in B (blue) neighborhoods. On the contrary, we find a 4.5 percentage point increase in the best-rated areas (A, green), but the caveats we mentioned in section 4.4 apply in this case.

In terms of African American percentage, we find a 1.8 percentage point increase in the lowest-rated areas (D, red) a 9.4% increase with respect to the baseline period, as reported in Table 8. We do not find any other effect of this policy in other areas, given the near absence of Black Americans in A, B and C neighborhoods. The results for the full set of specifications for this outcome can be found in Appendix Table A8. Table 8 also shows the estimated coefficients for property values. The assignment of C (yellow) and D (red) grades caused sizable reductions in property prices. While the reduction in property prices in red areas confirms the popular narrative for “redlining”, the negative effect for C neighborhoods is more surprising, and it was first documented in Aaronson et al. (2021b). Our empirical design does not find any significant effect in B areas.⁵⁶ Appendix Table A10 confirms the negative

55. In our research design, the treatment is assigned at the city level. Since in our data different periods are separated by 10-year gaps, we do not allow for within city serial correlation of standard errors. This clustering choice does not address unobserved, within-city, serially correlated shocks over a ten-year time span. One more threat is that hexagon-specific error components could be serially correlated across decades. Note that our results are robust to the inclusion of city fixed effects, which will absorb constant city-level error components, or neighborhood (hexagon) fixed effects. We report results with standard errors clustered at the city level, our most conservative option, in Appendix Table A15. The median ratio between city-year clustered standard errors and city clustered standard errors is 0.702, the average one is 0.782.

56. The results for the full set of specifications for property values can be found in Appendix Table A9. We are hesitant in interpreting the estimated coefficient β^A as the causal effect of grade A. The significantly smaller sample size, paired with weak evidence of parallel pre-trend for this class, invite caution when

effects in C and D areas when we apply a log transformation to property values, although the estimates are less precise. The discrepancy between the linear and logarithmic specifications suggests the presence of heterogeneous effects across the property value distribution. Since the logarithmic form reduces the contribution of higher prices, it seems that HOLC maps in 1940 had a stronger impact on more expensive houses in C and D areas relatively to lower-value ones. The patterns we find for property values do not translate to rent prices. The last column of Table 8 shows that the HOLC maps did not affect rental prices in 1940 for our cities of interest.

5.1.1 Robustness

All the short-term results are confirmed if we replicate our estimates using a neighborhood (hexagon) level dataset instead of an individual one. The coefficients and their standard errors can be found in Tables A11 to A14. These alternative specifications are estimated on a two-period panel of neighborhoods in our cities of interests. The stability of the results across the individual and the neighborhood levels datasets mitigates the concern that the short-term results might be driven by strong changes in cohort composition between 1930 and 1940.

In the results we discussed so far, observations were grouped according to the grade predicted by the random forest algorithm. In Table A16 we show that the results are robust when we replace predicted grades with observed HOLC classes for treated observations. Another robustness check is presented in Table A17 where we confirm that the estimated coefficients do not change if we extend the treatment group to cities up to 60,000 residents according to the 1930 census.⁵⁷

Our strategy might be capturing structural differences in the evolution of smaller versus bigger cities between 1930 and 1940. Table A18 shows that we do not find meaningful effects if we focus on placebo outcomes such as female percentage, number of children or male unemployment rate. As an additional check, we replicate our procedure with a placebo

considering the results for this grade.

57. Similar results can be obtained by changing the treatment-group population limit to 70,000, 80,000 and 100,000.

population threshold set at 60,000 people. This new threshold defines a new treatment group (cities between 60,000 and 70,000 residents) and a new control group (cities between 50,000 and 60,000 residents). Our research design should not replicate our main results with the new thresholds since HOLC practices did not differ between these two new sets of cities. Table A19 contains the results of this exercise. We do not find any relevant and significant effect for homeownership rates, African American percentage, and property values.⁵⁸

In section 4.4 we discussed how treatment spillovers from surrounding areas could be a threat to the validity of our analysis. As a first step to address this concern, Table A20 shows that our results are robust when we control for local grade composition. In particular, we include the prevalence of the four grades in a 1000mt (0.63miles) radius.

5.2 Long Term Results

At the time of writing, full individual census data are not available starting in 1950. To investigate the effects of HOLC maps in the second half of the twentieth century, we must employ alternative data sources. As mentioned in section 3.2.2, we use tract-level NHGIS data between 1960 and 2010. Unfortunately, this data source does not provide sufficient coverage for our cities of interest in 1950, so we drop this decade in the analysis.⁵⁹ We estimate the model described in equation (2) separately for each grade with neighborhoods (hexagons) as the unit of observations. Standard errors are clustered at the city-year level. Since census tracts are always bigger than hexagon neighborhoods in our cities of interest, the geographical variation underlying these estimates is much smaller than the variation we exploited in previous estimates.

Figure 11 shows the estimated DiD coefficients for homeownership rates and African American shares in C and D areas. We find reductions in homeownership rates between 4.4 and 5.6 percentage points in D (red) neighborhoods, while no statistically significant effects

58. The positive effect on property values in D areas has the opposite sign of what we find in our main results.

59. We plan to extend the current analysis to 1950 when full-count Census data for that year becomes available in April 2022.

can be found in C (yellow) ones. The increase in the local shares of Black Americans in D (red) areas we found in 1940 is confirmed in later years, as depicted in Figure 11. We also detect similar significant increases in C (yellow) zones. While the estimated effects in red areas are coherent with the short term results, the ones in yellow areas are more surprising. They show an increase between 5.9 and 7.9 percentage points in the percentage of African Americans in C areas starting in 1980. The complete results for these two outcomes can be found in Table 9. Because of their high level of geographical aggregation, NHGIS data do not provide enough information to characterize the long term effects of HOLC maps on property values. To make up for the lack of precision in these estimates, we include an additional data source in our long-term analysis of property values.

We turn to an alternative, more granular, source of information: the CoreLogic deeds. This additional dataset allows us to assess the impact of HOLC maps on real estate transactions between 1965 and 2005.⁶⁰ We estimate equation (2) with neighborhood level observations grouped in time bins with a 5-year frequency, so that the results are directly comparable with estimates obtained with NHGIS data. Each individual transaction is assigned a neighborhood, and hence a HOLC grade, using their geographic coordinates, available in CoreLogic. As in previous specifications, standard errors are clustered at the city-year level. Figure 12 shows the resulting coefficients and 95% confidence intervals. We find negative causal effects of HOLC maps on house values between 1965 and 1980 in D (red) and C (yellow) areas. The results for D neighborhoods describe a somewhat steady convergence of property values between treated and control cities in neighborhoods classified as D by the random forest algorithm. Statistically significant effects cannot be detected starting in mid 1980s. Table 10 shows the results for the four different HOLC grades.⁶¹ As we mentioned in section 2, between 1974 and 1977, three critical legislative measures were introduced with the primary goal of counteracting redlining in the mortgage market. Our results suggest that the combined

60. The choice of this time span is based on the coverage of CoreLogic deeds for our cities of interest. In particular, the dataset does not provide enough transactions to obtain reliable estimates prior to 1965. Additional details about CoreLogic’s coverage of our cities of interest can be found in Appendix Table A2. We stop the analysis in 2005 to avoid including the effects of the subprime mortgage crisis of the early 2000s.

61. The results are robust to a log-transformation of the outcome, as it is shown in Appendix Table A21 and Appendix Figure A6.

effects of the Equal Credit Opportunity Act, the Home Mortgage Disclosure Act and the Community Reinvestment Act might have been sufficient to offset persistent effects of HOLC maps on property prices.

6. Conclusion

In the second half of the 1930s, a federal agency undertook an unprecedented survey of the demographic and housing conditions of US neighborhoods in more than 200 cities. Its goal was to provide unified standards to assess real estate properties and stabilize a market that had just begun to recover from the Great Depression. The initiative was a data-driven effort based on the most advanced theories of urban development of the time, and its resulting maps were a data analytics tool in high demand among real estate professionals. Less than a hundred years later, the Home Owners' Loan Corporation maps have become a symbol of structural racism in the popular press and the political debate. Today's negative judgments of HOLC practices are based on non-discrimination principles that have guided US public institutions since the civil rights movement. Such condemnations are backed by historical evidence and are valid independent of quantitative estimates of causal effects. At the same time, measuring the consequences of the HOLC maps is an interesting exercise to understand the role of public institutions in coordinating and standardizing individual discriminatory behaviors.

The main challenge in estimating the causal effects of different HOLC grades is that the agency's personnel traced neighborhood borders and assigned evaluations with precision. Different HOLC grades within a city closely mirror socioeconomic trends we can observe in the census data. Instead of relying on spatial discontinuity designs, we take advantage of an exogenous threshold determining which cities the agency surveyed. Since we are interested in estimating the effects of different grades, we compare neighborhoods evaluated by HOLC with analogous neighborhoods in control cities. To classify neighborhoods in control cities, we train a random forest algorithm to replicate HOLC grades. Our spatial classification model has an out-of-sample accuracy of more than 90% and returns predicted maps that are

credible replicas of HOLC maps. Using the predicted grades, we then estimate the HOLC grades' short- and long-term effects with a grouped difference-in-differences design.

We find that this government intervention had adverse effects in areas that received the lowest grade. In the short-term, we find a 2.4 percentage points reduction in homeownership rates, a sizable reduction in property prices, together with a 1.8 percentage points increase in the percentage of African American residents. No effect on rent prices is found. We also find a 9.5% decrease in property values in C neighborhoods compared to the 1930 values. This consequence of “yellow-lining” is rarely discussed and was first highlighted by [Aaronson et al. \(2021b\)](#).

We have evidence that the negative effects of D grades in terms of homeownership rates and percentages of African American residents have persisted until the present, but we sometimes lack precision since the data are available only at the census tract level. For property prices, we exploit a more granular data source – CoreLogic deeds – to estimate the long-term evolution of the causal effects. We find significant negative effects on property prices in C and D neighborhoods until the early 1980s. This result differs from that of [Aaronson et al. \(2021b\)](#), who find significant effects on property prices up until 2010, using a different set of cities and an alternative identification strategy. In our case, the effects of the maps can no longer be detected in the decades following the introduction of legislation targeting residential redlining.

We have analyzed a government policy that institutionalized discrimination by standardizing appraisal standards. Given that discriminatory practices in the housing market were widespread at the time, it is not obvious that a graphical representation of mortgage risk could have affected homeownership rates or property prices. Our results show that an organization's acceptance and reproduction of discriminatory practices can have an economic effect. The HOLC maps could have influenced discrimination in the housing market via at least two mechanisms that are not mutually exclusive: first, replacing heterogeneous individual biases with one homogeneous set of biases; and second, solving an information asymmetry by circulating an information tool useful to discriminate between different neighborhoods. While our empirical strategy cannot differentiate between these two mechanisms, future research

could attempt to disentangle the roles of bias standardization and information provision to provide a more nuanced understanding of the consequences of institutional discrimination.

References

- Aaronson, Daniel, Jacob Faber, Daniel Hartley, Bhashkar Mazumder, and Patrick Sharkey (2021a) “The long-run effects of the 1930s HOLC “redlining” maps on place-based measures of economic opportunity and socioeconomic success,” *Regional Science and Urban Economics*, 86, 103622, <https://doi.org/10.1016/j.regsciurbeco.2020.103622>.
- Aaronson, Daniel, Daniel Hartley, and Bhashkar Mazumder (2021b) “The Effects of the 1930s HOLC “Redlining” Maps,” *American Economic Journal: Economic Policy*, 13 (4), 355–92, [10.1257/pol.20190414](https://doi.org/10.1257/pol.20190414).
- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010) “Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program,” *Journal of the American statistical Association*, 105 (490), 493–505.
- Abadie, Alberto and Javier Gardeazabal (2003) “The economic costs of conflict: A case study of the Basque Country,” *American economic review*, 93 (1), 113–132.
- Anders, John Paul (2019) “The Long Run Effects of Transformational Federal Policies: Redlining, the Affordable Care Act and Head Start.”
- Appel, Ian and Jordan Nickerson (2016) “Pockets of poverty: The long-term effects of redlining,” *Available at SSRN 2852856*.
- Arrow, Kenneth J. (1973) *The Theory of Discrimination*, 3–33: Princeton University Press, <http://www.jstor.org/stable/j.ctt13x10hs.5>.
- Becker, Gary S et al. (1971) “The Economics of Discrimination,” *University of Chicago Press Economics Books*.
- Bertrand, Marianne, Dolly Chugh, and Sendhil Mullainathan (2005) “Implicit Discrimination,” *American Economic Review*, 95 (2), 94–98, [10.1257/000282805774670365](https://doi.org/10.1257/000282805774670365).
- Bertrand, Marianne and Esther Duflo (2017) “Field experiments on discrimination,” *Handbook of economic field experiments*, 1, 309–393.
- Birch, Colin P.D., Sander P. Oom, and Jonathan A. Beecham (2007) “Rectangular and hexagonal grids used for observation, experiment and simulation in ecology,” *Ecological Modelling*, 206 (3), 347–359, <https://doi.org/10.1016/j.ecolmodel.2007.03.041>.
- Borusyak, Kirill and Peter Hull (2020) “Non-Random Exposure to Exogenous Shocks: Theory and Applications,” Working Paper 27845, National Bureau of Economic Research, [10.3386/w27845](https://doi.org/10.3386/w27845).
- Boustan, Leah Platt (2010) “Was postwar suburbanization “white flight”? Evidence from the black migration,” *The Quarterly Journal of Economics*, 125 (1), 417–443.
- (2013) “Racial Residential Segregation in American Cities,” Working Paper 19045, National Bureau of Economic Research, [10.3386/w19045](https://doi.org/10.3386/w19045).
- (2016) *Competition in the Promised Land*: Princeton University Press.

- Breiman, Leo (2001) “Random forests,” *Machine learning*, 45 (1), 5–32.
- Chyn, Eric (2018) “Moved to opportunity: The long-run effects of public housing demolition on children,” *American Economic Review*, 108 (10), 3028–56.
- Coates, Ta-Nehisi (2014) “The Case for Reparations,” *Atlantic*.
- Collins, William J and Katharine L Shester (2013) “Slum clearance and urban renewal in the United States,” *American Economic Journal: Applied Economics*, 5 (1), 239–73.
- Cutler, David M, Edward L Glaeser, and Jacob L Vigdor (1999) “The rise and decline of the American ghetto,” *Journal of Political Economy*, 107 (3), 455–506.
- Faber, Jacob W. (2020) “We Built This: Consequences of New Deal Era Intervention in America’s Racial Geography,” *American Sociological Review*, 85 (5), 739–775, [10.1177/0003122420948464](https://doi.org/10.1177/0003122420948464).
- Fang, Hanming and Andrea Moro (2011) “Theories of statistical discrimination and affirmative action: A survey,” *Handbook of social economics*, 1, 133–200.
- Federal Home Loan Bank Board (1937) “Fifth Annual Report of the Federal Home Loan Bank Board, For the Period July 1, 1936 through June 30, 1937.”
- Fishback, Price V, Alfonso Flores-Lagunes, William C. Horrace, Shawn Kantor, and Jaret Treber (2011) “The influence of the Home Owners’ loan corporation on housing markets during the 1930s,” *The Review of Financial Studies*, 24 (6), 1782–1813.
- Fishback, Price V, Jessica LaVoice, Allison Shertzer, and Randall Walsh (2020) “The HOLC Maps: How Race and Poverty Influenced Real Estate Professionals’s Evaluation of Lending Risk in the 1930s,” Working Paper 28146, National Bureau of Economic Research, [10.3386/w28146](https://doi.org/10.3386/w28146).
- Fishback, Price V, Jonathan Rose, Kenneth A Snowden, and Thomas Storrs (2021) “New Evidence on Redlining by Federal Housing Programs in the 1930s,” Technical report, National Bureau of Economic Research.
- Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther (2021) “Predictably Unequal? The Effects of Machine Learning on Credit Markets,” *Journal of Finance*, *Forthcoming*.
- Glaeser, Edward and Jacob Vigdor (2012) “The end of the segregated century,” *Manhattan Institute for Policy Research*, *January*, 23–26.
- Hillier, Amy E (2003) “Redlining and the Home Owners’ Loan Corporation,” *Journal of Urban History*, 29 (4), 394–420.
- (2005) “Residential security maps and neighborhood appraisals: the Home Owners’ Loan Corporation and the case of Philadelphia,” *Social Science History*, 29 (2), 207–233.
- Ihlanfeldt, Keith R and Benjamin Scafidi (2002) “Black self-segregation as a cause of housing segregation: Evidence from the multi-city study of urban inequality,” *Journal of Urban Economics*, 51 (2), 366–390.
- Jackson, Kenneth T (1980) “Race, ethnicity, and real estate appraisal: The Home Owners Loan Corporation and the federal housing administration,” *Journal of Urban History*, 6 (4), 419–452.

- Jones-Correa, Michael (2000) “The origins and diffusion of racial restrictive covenants,” *Political Science Quarterly*, 115 (4), 541–568.
- Krimmel, Jacob (2018) “Persistence of prejudice: Estimating the long term effects of redlining.”
- Krysan, Maria and Reynolds Farley (2002) “The residential preferences of blacks: Do they explain persistent segregation?” *Social Forces*, 80 (3), 937–980.
- Lang, Kevin and Ariella Kahn-Lang Spitzer (2020) “Race Discrimination: An Economic Perspective,” *Journal of Economic Perspectives*, 34 (2), 68–89, [10.1257/jep.34.2.68](https://doi.org/10.1257/jep.34.2.68).
- Lieberman, Andres, Christopher Neilson, Luis Opazo, and Seth Zimmerman (2018) “The Equilibrium Effects of Information Deletion: Evidence from Consumer Credit Markets,” Working Paper 25097, National Bureau of Economic Research, [10.3386/w25097](https://doi.org/10.3386/w25097).
- Light, Jennifer S (2010) “Nationality and neighborhood risk at the origins of FHA underwriting,” *Journal of Urban History*, 36 (5), 634–671.
- Logan, John R and Weiwei Zhang (2018) “Developing GIS Maps for US Cities in 1930 and 1940,” in *The Routledge Companion to Spatial History*, 229–249: Routledge.
- Logan, Trevon D. and John M. Parman (2017) “The National Rise in Residential Segregation,” *The Journal of Economic History*, 77 (1), [10.1017/S0022050717000079](https://doi.org/10.1017/S0022050717000079).
- Ludwig, Jens and Sendhil Mullainathan (2021) “Fragile Algorithms and Fallible Decision-Makers: Lessons from the Justice System,” Working Paper 29267, National Bureau of Economic Research, [10.3386/w29267](https://doi.org/10.3386/w29267).
- Manson, Steven, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles (2021) “IPUMS National Historical Geographic Information System: Version 16.0,” <http://doi.org/10.18128/D050.V16.0>.
- Michney, Todd M. (2021) “How the City Survey’s Redlining Maps Were Made: A Closer Look at HOLC’s Mortgagee Rehabilitation Division,” *Journal of Planning History*, 0 (0), 15385132211013361, [10.1177/15385132211013361](https://doi.org/10.1177/15385132211013361).
- Mullainathan, Sendhil and Jann Spiess (2017) “Machine learning: an applied econometric approach,” *Journal of Economic Perspectives*, 31 (2), 87–106.
- Nagaraj, Abhishek and Scott Stern (2020) “The Economics of Maps,” *Journal of Economic Perspectives*, 34 (1), 196–221, [10.1257/jep.34.1.196](https://doi.org/10.1257/jep.34.1.196).
- Nelson, Robert K., LaDale Winling, Richard Marciano, and Nathan Connolly (2021) “Mapping Inequality.”
- Phelps, Edmund S (1972) “The statistical theory of racism and sexism,” *The American Economic Review*, 62 (4), 659–661.
- Probst, Philipp, Marvin Wright, and Anne-Laure Boulesteix (2018) “Hyperparameters and Tuning Strategies for Random Forest,” *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, [10.1002/widm.1301](https://doi.org/10.1002/widm.1301).

- Rambachan, Ashesh, Jon Kleinberg, Jens Ludwig, and Sendhil Mullainathan (2020) “An economic perspective on algorithmic fairness,” in *AEA Papers and Proceedings*, 110, 91–95.
- Rothstein, Richard (2017) *The color of law: A forgotten history of how our government segregated America*: Liveright Publishing.
- Ruggles, Steven, Catherine A. Fitch, Ronald Goeken et al. (2020) “IPUMS Restricted Complete Count Data: Version 2.0.”
- Small, Mario L. and Devah Pager (2020) “Sociological Perspectives on Racial Discrimination,” *Journal of Economic Perspectives*, 34 (2), 49–67, [10.1257/jep.34.2.49](https://doi.org/10.1257/jep.34.2.49).
- Sood, Aradhya, William Speagle, and Kevin Ehrman-Solberg (2019) “Long Shadow of Racial Discrimination: Evidence from Housing Covenants of Minneapolis.”
- Tach, Laura and Allison Dwyer Emory (2017) “Public housing redevelopment, neighborhood change, and the restructuring of urban inequality,” *American journal of sociology*, 123 (3), 686–739.

Table 1 — 1930 Population Distribution According to HOLC Grades

	N	Population Shares			
		A	B	C	D
General Population	30,945,584	3.9%	18.3%	41.7%	36.1%
By Race					
White	28,801,136	4.1%	19.4%	43.9%	32.6%
Black	2,094,493	0.9%	3.1%	12.0%	83.6%

Notes: The sample includes 1930 Census individuals with a valid geocode in neighborhoods with a digitized HOLC map.

Table 2 — 1930 Descriptive Statistics According to HOLC Grades

	Grade			
	A	B	C	D
Black	0.01 (0.09)	0.01 (0.085)	0.02 (0.11)	0.17 (0.32)
Home Owner	0.76 (0.28)	0.66 (0.27)	0.56 (0.27)	0.43 (0.28)
House Value	12,938 (7,013)	8,805 (5,134)	6,638 (3,967)	5,038 (3,648)
Rent	143 (352)	73 (174)	51 (120)	42 (112)
Income Score	7.30 (0.26)	7.23 (0.22)	7.12 (0.21)	6.931 (0.30)
First Gen Immigrant	0.14 (0.17)	0.16 (0.17)	0.22 (0.20)	0.23 (0.24)
Unemployed, Men	0.04 (0.11)	0.06 (0.11)	0.10 (0.13)	0.13 (0.14)
Owns a Radio	0.76 (0.27)	0.69 (0.26)	0.58 (0.27)	0.38 (0.27)

Notes: The sample includes 1930 census individuals with a valid geocode in neighborhoods with a digitized HOLC map. See Appendix Section A.2 for definitions of Census variables in our dataset.

Table 3— Random Forest Performance, Confusion Matrix

		Data			
		D	C	B	A
Prediction	D	12940	668	62	3
	C	927	18771	827	74
	B	93	792	9939	462
	A	7	36	153	2792
Accuracy		91.55%			
Class Sensitivity		92.65	92.62	90.51	83.82
Prevalence		28.77	41.75	22.62	6.86
Detection Prevalence		28.17	42.43	23.25	6.15

Notes: The matrix compares the observed and the predicted grades for a test set of observations excluded from the training procedure. The test set is a 25% random subsample of the original dataset selected with stratified sampling according to city population and HOLC grade. The level of observation is a neighborhood (hexagon). See Section 3.1 for details about the hexagon definition. The sample includes every hexagon in a mapped city with a 1930 population below 3,000,000 and containing at least 20 residents in 1930. The results are robust to different sample definitions in terms of population cutoffs. A predicted grade is the class predicted by the trained random forest algorithm. See Section 4.1 and Appendix Section A.3 for details about the Random Forest training procedure. Overall Accuracy is the percentage of hexagons whose predicted grades correspond to observed ones. Class sensitivity for a grade j is the proportion of correctly predicted hexagons among the spatial units with grade j . Prevalence reports the share of each observed grade in the test set, while detection prevalence shows the distribution of predicted grades.

Table 4— Shares of Local Grades, by Grade and Treatment Status

	A	B	C	D
Share A	50.2%	4.9%	0.5%	0.3%
Share B	38.9%	64.0%	11.3%	4.6%
Share C	9.3%	27.4%	76.7%	32.3%
Share D	1.6%	3.6%	11.5%	62.8%

Notes: The Table reports average shares of surrounding grades according to neighborhoods grades and treatment status. Neighborhood surroundings are defined with a 1000mt. radius (0.63 miles). The sample includes neighborhoods with at least 20 residents in 1930 in cities with a population between 30,000 and 50,000. See Appendix Section A.4 for a list of cities.

Table 5 — Testing Differences in 1930-1920 Trends by Treatment Status

Testing 1930-1920 Trends By Treatment Status

<i>Dependent Variable</i>	A	B	C	D
Black	0.008 (0.006)	0.001 (0.002)	-0.001 (0.003)	-0.003 (0.010)
Home Owner	-0.080** (0.034)	-0.017 (0.019)	-0.006 (0.015)	-0.021 (0.014)
Income Score	-0.023 (0.026)	-0.019 (0.012)	-0.006 (0.008)	0.021* (0.012)
Education Score	-2.055 (1.413)	-0.478 (0.554)	0.173 (0.259)	0.074 (0.289)
First Gen. Immigrant	0.018 (0.016)	0.004 (0.007)	0.014** (0.007)	-0.012 (0.012)
Number of Children	0.014 (0.021)	0.004 (0.011)	0.011 (0.008)	0.001 (0.010)

Notes: The Table reports the coefficients from a set of regressions where the dependent variable is the 1930-1920 change in the variable reported in the left column, and the independent variable is an indicator for treatment status. See Appendix Section A.2 for definitions of Census variables in our dataset. The level of observation is a neighborhood (hexagon). The sample includes every hexagon in cities with a 1930 population between 30,000 and 50,000 and at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table 6 — 1930 Descriptive Statistics According to HOLC and Predicted Grades

	Grade			
	C		D	
	HOLC	Predicted	HOLC	Predicted
Black	0.02 (0.10)	0.02 (0.08)	0.17 (0.31)	0.21 (0.33)
Home Owner	0.53 (0.23)	0.52 (0.23)	0.41 (0.24)	0.40 (0.23)
Property Value	6,843 (3,841)	6,869 (3,861)	5,253 (3,632)	4,729 (3,267)
Rent	52 (112)	54 (116)	42 (98)	39 (97)
Income Score	7.13 (0.18)	7.13 (0.17)	6.91 (0.28)	6.86 (0.28)
First Gen Immigrant	0.22 (0.17)	0.23 (0.18)	0.24 (0.22)	0.23 (0.23)
Unemployed, Men	0.10 (0.09)	0.10 (0.08)	0.13 (0.10)	0.13 (0.10)
Owns a Radio	0.58 (0.22)	0.58 (0.21)	0.38 (0.23)	0.33 (0.21)

Notes: The Table reports averages of 1930 census variables according to different classifications. The first two columns compare means between hexagons classified as *C* by the HOLC with those classified as *C* by our random forest algorithm. The third and fourth columns do the same for grade *D*. The level of observation is a neighborhood (hexagon). The sample includes all the hexagons intersecting a HOLC neighborhood digitized by Nelson et al. (2021) in 202 maps. Standard deviations are reported in parentheses.

Table 7 — Short-Term Difference-in-Differences Results, by Grade

	<i>Dependent variable: Homeownership Rates</i>			
	(1)	(2)	(3)	(4)
<i>DiD_A</i> $\bar{Y}^A = 0.61$	0.032 (0.093)	0.047*** (0.016)	0.031** (0.014)	0.045*** (0.015)
<i>DiD_B</i> $\bar{Y}^B = 0.62$	0.002 (0.027)	0.003 (0.009)	0.007 (0.009)	-0.002 (0.009)
<i>DiD_C</i> $\bar{Y}^C = 0.49$	-0.010 (0.031)	-0.012 (0.007)	-0.009 (0.007)	-0.017** (0.007)
<i>DiD_D</i> $\bar{Y}^D = 0.39$	-0.022 (0.044)	-0.017** (0.008)	-0.015** (0.008)	-0.024*** (0.009)
City Fixed. Eff.		X		X
Hexagon Fixed Eff.			X	
Hexagon Controls				X

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade. Each row contains the DiD coefficients for a given grade. \bar{Y}^J is the average for the outcome of interest, homeownership rates, in the pre-treatment period (1930) for grade J . The first column report the DiD coefficients resulting from a simple DiD framework. The second one replaces the indicator for treatment, which is assigned at the city level, with a city fixed effect, while the third replaces it with a neighborhood (hexagon) fixed effect. Column 4 reports the DiD coefficients when we add geographic and demographic controls at the hexagon level to the specification from the second column. The list of controls includes geographic coordinates, a scaled measure of distance from the city center, spatial unit's population density, imputed income score and family size. The regressions are estimated with individual-level observations. The sample includes individuals with valid geocodes in cities with a 1930 population between 30,000 and 50,000, living in hexagons with at least 20 residents in 1930. See Appendix Section A.3 for a list of cities. The number of observations according to each grade, N^J , are: $N^A = 137,144$, $N^B = 979,145$, $N^C = 3,116,521$, $N^D = 1,195,213$. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table 8 — Short-Term Difference-in-Differences Results, by Grade

	<i>Dependent Variables</i>		
	African American Percentage	Property Values	Rent Prices
DiD_A	-0.005 (0.006)	1,046*** (378)	39.3*** (12.0)
DiD_B	0.006* (0.004)	-106 (170)	8.8 (12.1)
DiD_C	0.001 (0.001)	-502*** (149)	5.9 (9.6)
DiD_D	0.018*** (0.004)	-302** (153)	3.1 (8.1)
\bar{Y}^A	0.008	9,305	107.4
\bar{Y}^B	0.011	6,836	57.0
\bar{Y}^C	0.020	5,274	43.4
\bar{Y}^D	0.190	3,500	29.0

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade for three different outcomes. Each row contains the DiD coefficients for a given grade. \bar{Y}^J is the average for the outcome of interest in the pre-treatment period (1930) for grade J . The Table shows the DiD coefficients resulting from a DiD framework with a city fixed effect and geographic and demographic controls at the hexagon level. The regression specification is analogous to the one in column (4) of Table 7. The list of controls includes geographic coordinates, a scaled measure of distance from the city center, spatial unit's population density, imputed income score and family size. The regressions are estimated with individual-level observations. The sample includes individuals with valid geocodes in cities with a 1930 population between 30,000 and 50,000, living in hexagons with at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table 9 — Long-Term Difference-in-Differences Results, by Grade. Census Data

	<i>Dependent variable:</i>			
	Homeownership Rates		African American Percentage	
	C	D	C	D
DiD_{60}	-0.003 (0.015)	-0.045** (0.019)	0.015 (0.017)	0.080*** (0.027)
DiD_{70}	-0.019 (0.014)	-0.047** (0.019)	0.036** (0.015)	0.077*** (0.023)
DiD_{80}	-0.015 (0.014)	-0.046** (0.020)	0.059*** (0.014)	0.098*** (0.025)
DiD_{90}	-0.010 (0.013)	-0.056*** (0.018)	0.072*** (0.016)	0.089*** (0.024)
DiD_{00}	-0.020 (0.012)	-0.054*** (0.017)	0.079*** (0.016)	0.082*** (0.025)
DiD_{10}	-0.026** (0.013)	-0.044** (0.019)	0.078*** (0.017)	0.068*** (0.025)
N	104,887	41,328	104,899	41,332
R^2	0.225	0.376	0.491	0.580

Notes: The Table reports Difference-in-Differences coefficients obtained estimating equation (2) by grade. Each row contains the DiD coefficients for a given grade in the corresponding year. In the reported specification, we replace the indicator for treatment, which is assigned at the city level, with a city fixed effect. The list of controls includes geographic coordinates and their squares together with state-specific linear time trends. The regressions are estimated with neighborhood (hexagon) level observations. Observations are weighted by log-transformed 1930 neighborhood population. The sample includes neighborhoods in cities with a 1930 population between 30,000 and 50,000, with at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. The data source for post-1940 outcomes is NHGIS; see Section 3.2.2 for details. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table 10 — Long-Term Difference-in-Differences Results, by Grade. CoreLogic

	<i>Dependent variable: Property Values</i>			
	Grade			
	A	B	C	D
DiD_{65}	. (.)	292 (3,781)	-12,360*** (2,689)	-13,669*** (1,976)
DiD_{70}	-10,514 (8,839)	-373 (4,510)	-16,607*** (3,547)	-17,601*** (2,790)
DiD_{75}	-7,939 (8,669)	-1,259 (4,449)	-13,155*** (3,087)	-14,949*** (2,804)
DiD_{80}	-8,876 (9,052)	-6,208 (5,168)	-21,296*** (6,009)	-11,320** (4,708)
DiD_{85}	18,516** (9,435)	12,686 (10,943)	6,112 (11,426)	-6,907 (7,885)
DiD_{90}	15,363 (10,427)	9,267 (7,779)	2,030 (8,101)	-5,321 (8,057)
DiD_{95}	-1,621 (11,155)	-904 (6,663)	-7,065 (5,933)	-9,326* (4,928)
DiD_{00}	6,795 (14,086)	3,297 (11,277)	2,978 (9,409)	-3,095 (7,059)
DiD_{05}	20,774 (17,199)	-5,353 (12,564)	5,457 (11,622)	-342 (9,584)
N	5,399	40,267	97,015	28,421
R^2	0.150	0.132	0.169	0.266

Notes: The Table reports Difference-in-Differences coefficients obtained estimating equation (2) by grade. Each row contains the DiD coefficients for a given grade in the corresponding year. The regressions are estimated with neighborhood-level observations. The sample includes neighborhoods with at least 20 residents in 1930 in cities with a population between 30,000 and 50,000. See Appendix Section A.4 for a list of cities. The data source for post-1940 outcomes is CoreLogic, see Section 3.3 for details. The outcome variable is adjusted with CPI to 1980 dollars. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

HOLC Residential Security Map of New Haven, CT

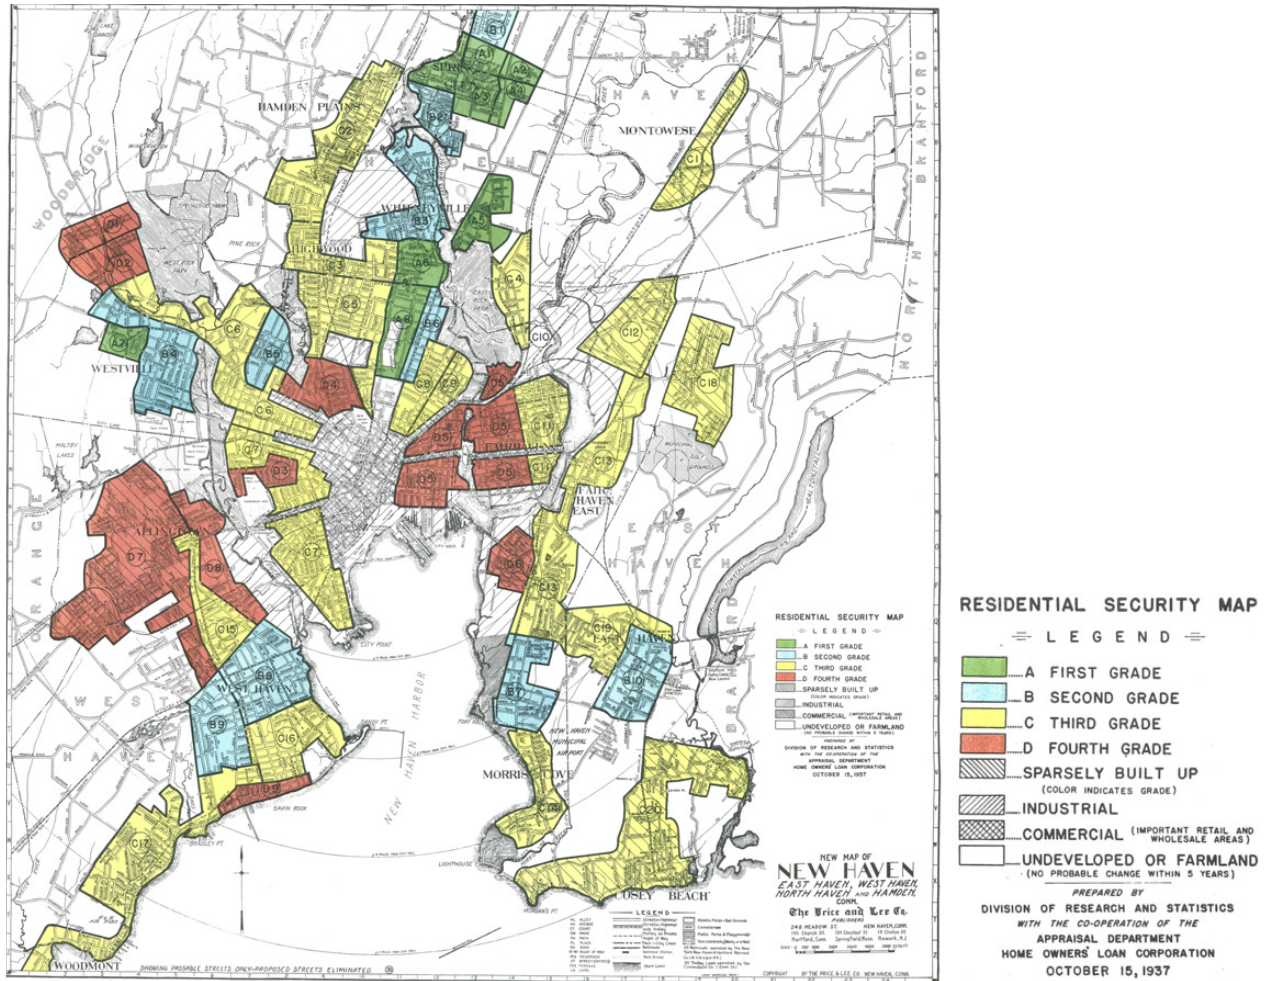


Figure 1 — The scan of the original Residential Security Map of New Haven, CT has been provided by *Mapping Inequality* (Nelson et al., 2021).

HOLC Area Description, Neighborhood D-4, New Haven, CT

NS FORM-8
8-28-37

AREA DESCRIPTION

1. NAME OF CITY NEW HAVEN, CONN. SECURITY GRADE FOURTH AREA NO. D-4

2. DESCRIPTION OF TERRAIN. Flat land with tree lined streets.

3. FAVORABLE INFLUENCES. Convenient to center of city.

4. DETRIMENTAL INFLUENCES. Age and obsolescence of dwellings as well as character of development and inhabitant.

5. INHABITANTS:

a. Type Domestics ; b. Estimated annual family income \$ 900.00

c. Foreign-born Mixed ; 50 % ; d. Negro Yes ; 70 % ;
(Nationality) (Yes or No)

e. Infiltration of Negro ; f. Relief families Many ;

g. Population is increasing ; decreasing ; static.

6. BUILDINGS:

a. Type or types 1, 2 & 3 family ; b. Type of construction Frame, few brick ;

c. Average age 25 to 75 years ; d. Repair Poor

7. HISTORY:

YEAR	SALE VALUES			RENTAL VALUES		
	RANGE	PREDOM- INATING	%	RANGE	PREDOM- INATING	%
1929 level	<u>\$5M - \$20M</u>	<u>8M</u>	<u>100%</u>	<u>\$12 1/2 - \$35</u>	<u>\$25</u>	<u>100%</u>
1935 low	<u>2,5M - 10M</u>	<u>4M</u>	<u>50%</u>	<u>9 - 22 1/2</u>	<u>17 1/2</u>	<u>70%</u>
1937 current	<u>2,5M - 10M</u>	<u>4M</u>	<u>50%</u>	<u>10 - 25</u>	<u>20</u>	<u>75%</u>

Peak sale values occurred in 1922 and were 100 % of the 1929 level.

Peak rental values occurred in 1929 and were 100 % of the 1929 level.

8. OCCUPANCY: a. Land 100 % ; b. Dwelling units 90 % ; c. Home owners 20 %

9. SALES DEMAND: a. None ; b. ; c. Activity is None

10. RENTAL DEMAND: a. Poor ; b. Units \$10 - \$25 ; c. Activity is Poor

11. NEW CONSTRUCTION: a. Types ; b. Amount last year None

12. AVAILABILITY OF MORTGAGE FUNDS: a. Home purchase None ; b. Home building None

13. TREND OF DESIRABILITY NEXT 10-15 YEARS Further downward

14. CLARIFYING REMARKS:
This is an older section of the city now given over largely to Negroes employed as domestics. Dwellings vary from small singles to multi-family. Section is quite congested and gives the appearance of a slum area. Absence of market has resulted in some demolition. Section is subject to vandalism.

15. Information for this form was obtained from See Explanations

Date October 15th 1937

Figure 2— The scan of the original Area Description for neighborhood D-4 of New Haven, CT has been provided by *Mapping Inequality* (Nelson et al., 2021).

Comparison of HOLC Digitized Map and its Hexagon Version

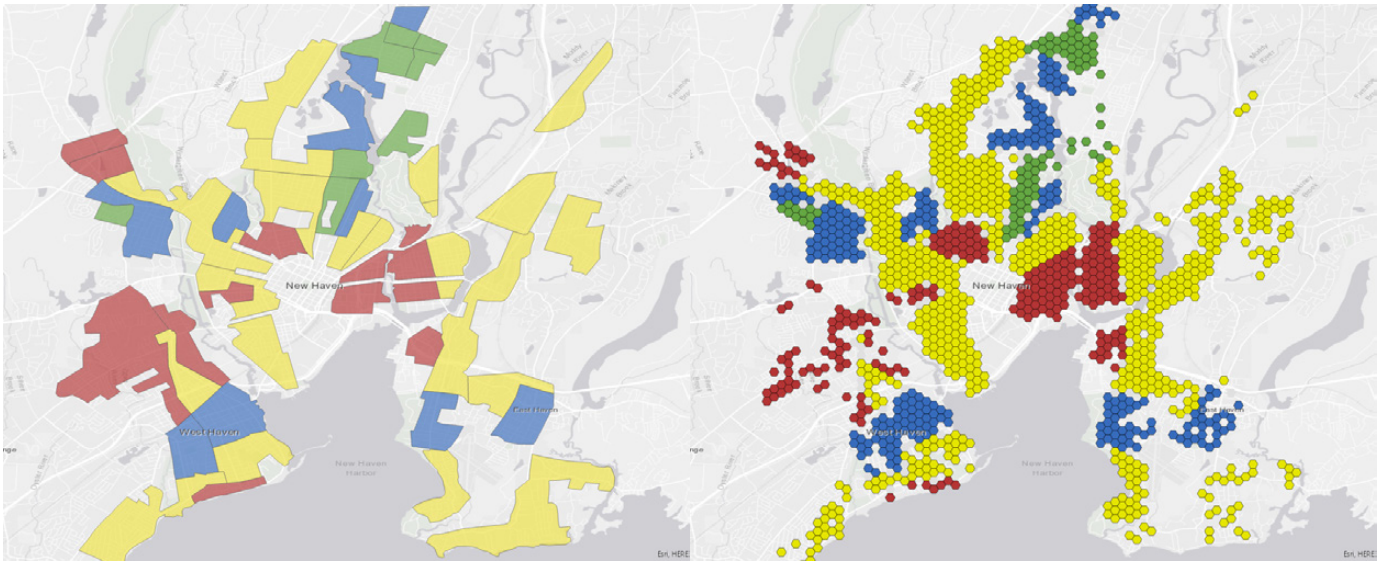


Figure 3— The digitized version of the Residential Security Map of New Haven, CT, shown in the left panel, has been provided by *Mapping Inequality* (Nelson et al., 2021). Details about the definition of the hexagon grid can be found in Section 3.1. The right panel shows our hexagon-level replica of the original HOLC map. All the maps are north-oriented.

Long Term Trends by HOLC Grade

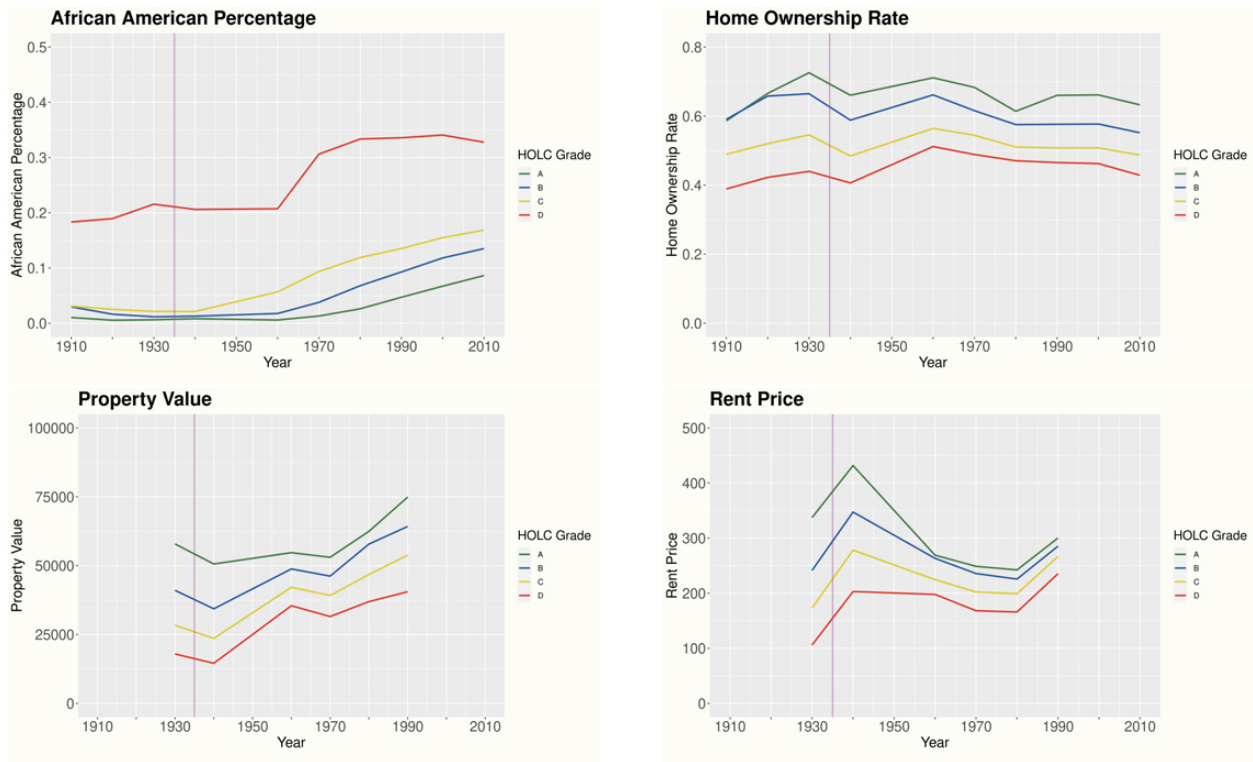


Figure 4— The sample includes neighborhoods located within a digitized HOLC map. See Appendix Section A.2 for definitions of Census variables in our dataset. The data sources are US full count census for 1910-1940 and NHGIS for 1960-2010.


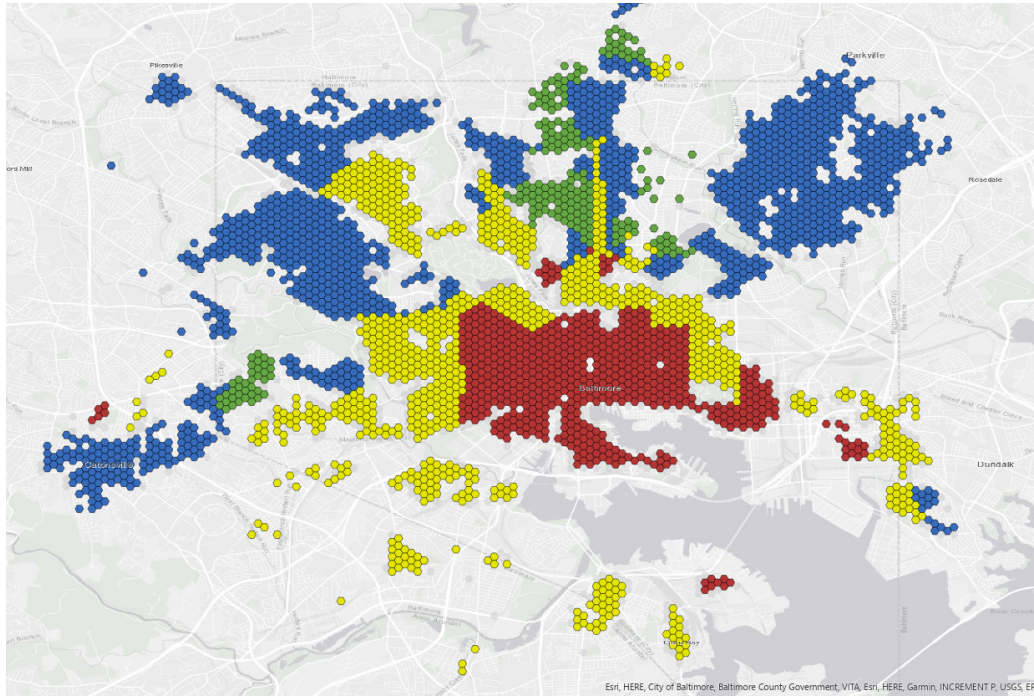
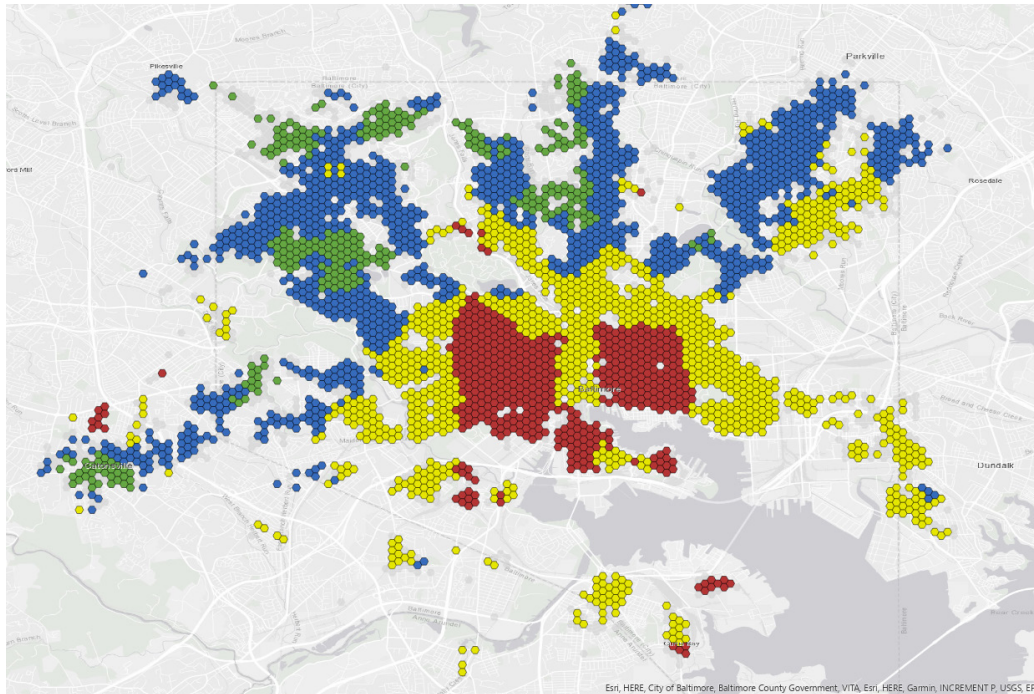


Figure 5 — The graph shows the treatment status of US cities according to their 1930 population. The vertical line highlights the 40,000 people threshold. Orange points identify cities in the control group, while the color purple highlights treated cities.

Comparison of HOLC and Predicted Maps for Baltimore, MD



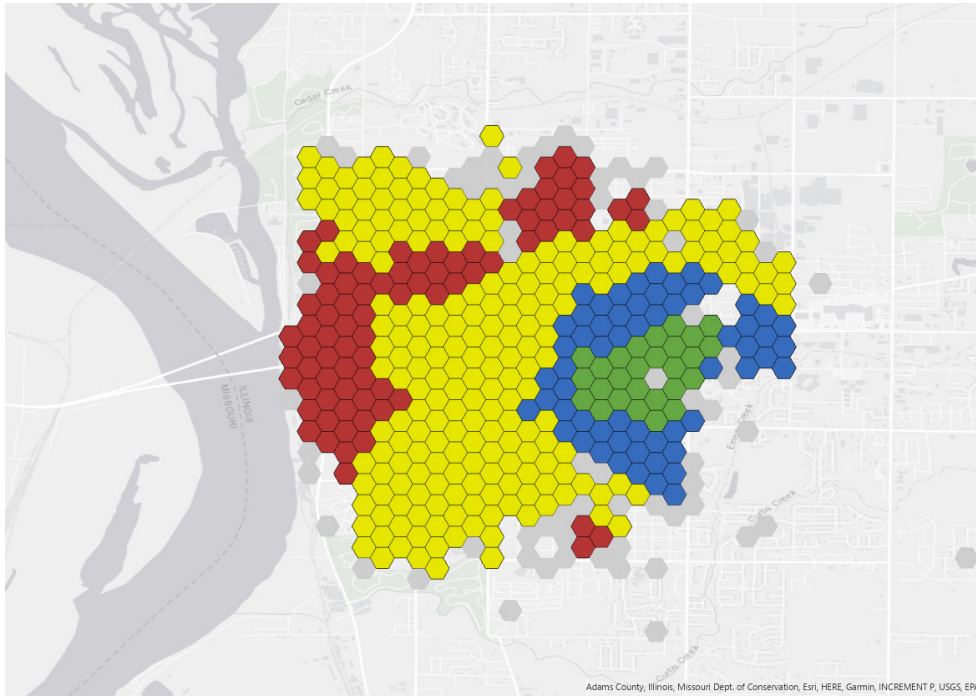
(a) HOLC Map (Nelson et al., 2021)



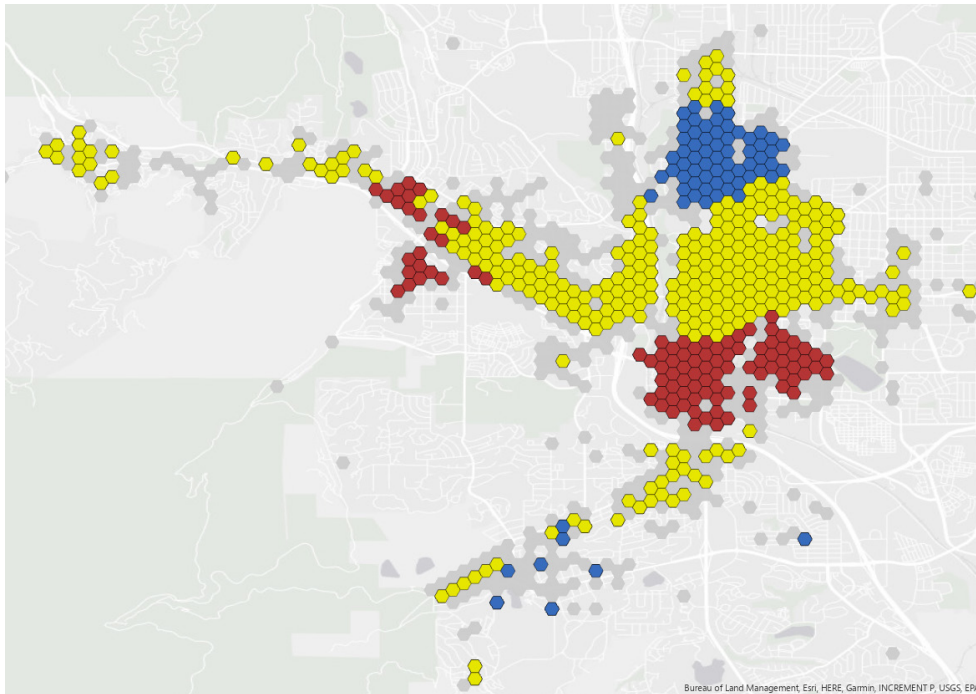
(b) Predicted Map, Random Forest Algorithm

Figure 6 — The Figure compares the digitized version of the HOLC maps for Baltimore, MD (Nelson et al., 2021) with the hexagon-level map we predict with the trained random forest algorithm. The correspondence between colors and grades is: Green=A, Blue= B, Yellow=C, Red=D. Grey hexagons have less than 20 residents in 1930 and are excluded from the prediction exercise. All the maps are north-oriented.

Predicted Maps in Control Cities



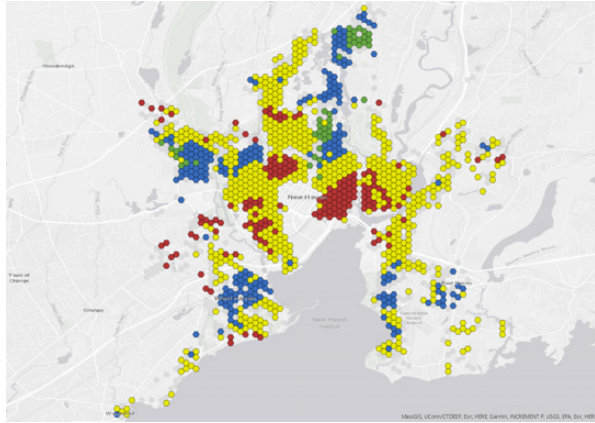
(a) Quincy, IL



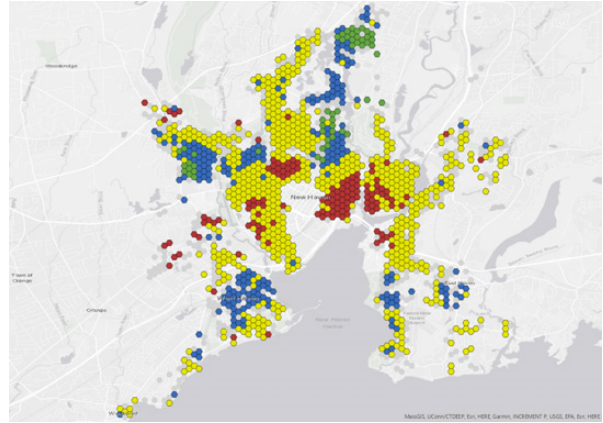
(b) Colorado Springs, CO

Figure 7 — The Figure compares the hexagon-level maps predicted with the trained random forest algorithm for Quincy, IL and Colorado Springs, CO. The correspondence between colors and grades is: Green=A, Blue= B, Yellow=C, Red=D. Grey hexagons have less than 20 residents in 1930 and are excluded from the prediction exercise. All the maps are north-oriented.

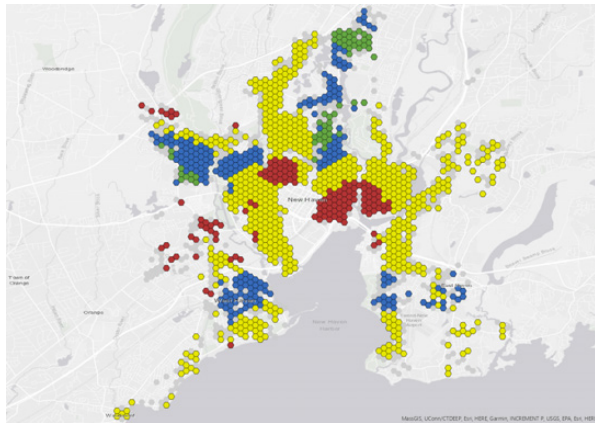
Comparison of Predicted Maps with Different Training Datasets



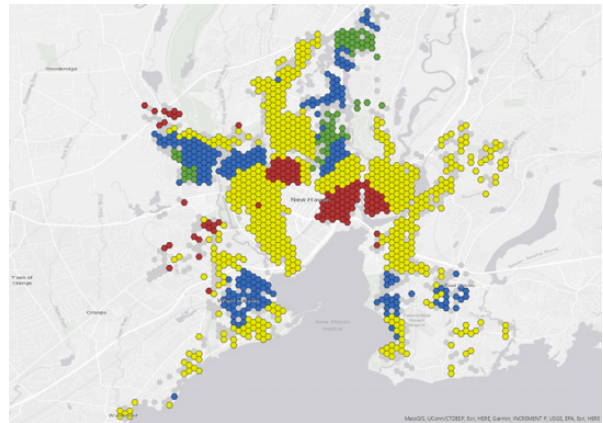
(a) Hexagon-Level



(b) Hexagon-Level, City-Level



(c) Hexagon-Level and Local Level



(d) Hexagon-Level, Local-Level, City-Level

Figure 8 — The maps show neighborhoods (hexagons) for New Haven, CT. The colors represent the grade predicted by the random forest algorithm. The correspondence between colors and grades is: Green=A, Blue= B, Yellow=C, Red=D. Grey hexagons have less than 20 residents in 1930 and are excluded from the prediction exercise. Different panels show predicted grades for random forests trained on four different sets of variables. The training datasets differ in terms of their levels of geographical aggregation, but not because of the variables included. The top-left panel shows predicted grades when only hexagon-level variables are included. The top-right panel adds city-level variables. The bottom-left panel replaces city-level variables with local-level information about the surrounding area. The surrounding area includes any hexagon whose centroids is within a 500mt. or 1000mt. radius. The bottom-right panel shows the predicted grades when we include all the previously mentioned variables. See Section 4.1 for details about the training of our random forest algorithm. All the maps are north-oriented.

Observable Pre-Trends, Grades C and D



Figure 10 — The Figure shows pre-trends for selected variables for C and D grades. The point estimates are averages of hexagon-level observations. The bars show the respective standard errors of each mean. The sample includes hexagons in cities with a 1930 population between 30,000 and 50,000, with at least 20 residents in 1930. The vertical line highlights 1930, the last pre-treatment decade. See Appendix Section A.2 for definitions of census variables in our dataset.

Long-Term Difference-in-Differences Results, C and D Grades. Census Data

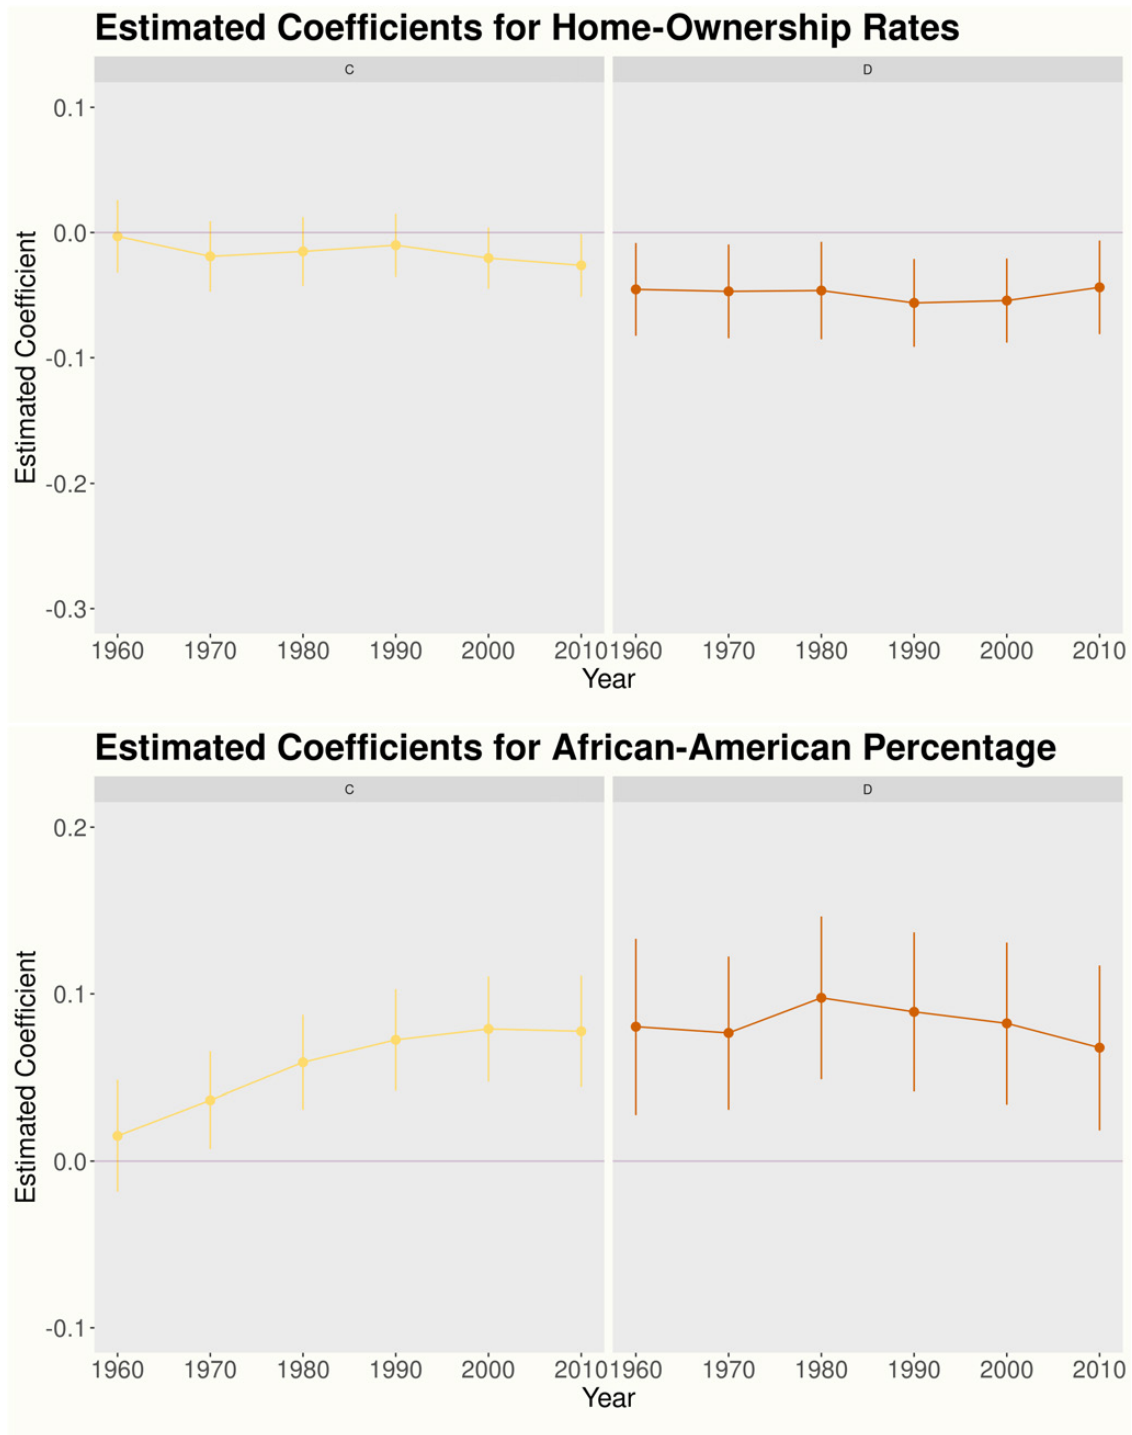


Figure 11 — The Figure shows the estimated coefficients for regression (2) and their 95% confidence intervals for homeownership rates and African American percentage. The Figure includes the results for grades C and D. The coefficients and standard errors are the ones reported in Table 9. See the Notes of Table 9 for estimation details.

Long-Term Difference-in-Differences Results, C and D Grades. CoreLogic Data




Figure 12 — The Figure shows the estimated coefficients for regression (2) and their 95% confidence intervals for property values. The Figure includes the results for grades C and D. The coefficients and standard errors are the ones reported in Table 10. See the Notes of Table 10 for estimation details.

A. Appendix

A.1 Geocoding Procedure

A.1.1 Address Cleaning

We clean addresses in decades between 1910 and 1940 following the procedure outlined in Logan and Zhang (2018). In particular:

- We cleaned street names. Names containing geographic indicators were removed, if they were not street names, and dummy variables were created for group quarters (hotels, apartments, convents, hospitals, group homes). House number information was extracted from street names.
- We cleaned house numbers. If the number found in the house number variable conflicted with the house number extracted from the street variable and the home was rented, the house number variable was interpreted to represent an apartment number.
- We interpolated missing street names and house numbers, conservatively. For observations on the same census page and within 6 house numbers from one another, missing streets were given the street name of the prior observation. For rented homes, missing house numbers were given the house number of the prior observation. For owned homes, if the street name was the same as the prior observation, missing house numbers were assigned a value equal to the house number of the prior observation plus two.

A.1.2 Geocoding

We geocoded the head of each household using ESRI Streetmap Premium 2019. These new-generation locator combines street addresses routing coordinates and parcel centroids databases to improve the number and the quality of the matches. Each address-coordinate match is assigned a 0-100 score by ESRI algorithm. We include in our analysis only matches with a score of at least 85. This choice is rather conservative and reduces measurement errors due to wrong locations of census households.

A.2 Census Variables Definitions

The following definitions are based on information provided by IPUMS.com documentation.

Census Variable Definitions	
Homeownership	Indicates whether the housing unit was owned, rather than rented, by its inhabitants.
African American	Based on census race variable. Prior to 1960, the census enumerator was responsible for categorizing persons and was not specifically instructed to ask the individual his or her race.
Property Values	For 1930 and 1940, enumerators consulted with the owners to estimate the sale value of the housing unit.
Rent Prices	Amount of the household's monthly contract rent payment.
First-Generation-Migrant	Whether an individual was foreign born. Based on the census variable <i>nativity</i> .
Unemployed, Men	Indicator defined according to census variable <i>empstat</i> for men between 18 and 65 years of age.
Radio Ownership	Whether any member of a family or housing unit owns a radio set.
Education Score	Census-built percentage of people in the respondent's occupational category who had completed one or more years of college
Number of Children	The number of own children residing with each individual.
Population Density	For any neighborhood the ratio between the area population and surface. Hexagon surface is fixed at $0.025km^2$.

A.2.1 Income Score Imputation

We estimate a log-wage regression on 1940 census data focusing on men aged 25-55 living in urban areas who were employed for wages. We regress self-reported wage income on a second degree polynomial in age, dummies indicating black, hispanic and immigration status, 3-digit occupation and state of residence indicators. Moreover, we include interactions between each of black, hispanic and immigration status with the age polynomial and interactions between each of the demographic variables with 1-digit occupations and state of residence.

The imputed income score in decades 1930, 1920, 1910 is the prediction based on the resulting estimates for men aged 25-55 who are employed for wages in those years.

A.3 Random Forest Training Procedure

We train the random forest algorithm with a hexagon-level dataset ($N = 192,016$) containing all cities mapped by the HOLC with a population up to 3,000,000. The dataset includes 48 different 1930 census variables. The variables are included at different geographical levels, bringing the total number of training variables to 163. The geographical levels employed in the training procedure are: hexagon, hexagon surroundings (500mt and 1000mt), city, county. In particular the variables are:

Random Forest Training Variables

- Share of African Americans
- Share of Women
- Share of Home-Owners
- Share of Population Not Speaking English
- Share Married
- Share of families owning a Radio
- Family Size
- Number of Children
- Age at First Marriage
- House Values
- Rent Prices
- Imputed Income
- Earning Scores
- Educational Scores
- House Distance from City Center
- Neighborhood Population Density
- Labor Force Participation, by gender
- Unemployment Rates, by gender
- Self-Employed and waged employees
- Share of First Generation Immigrants
- Share of Second Generation Immigrants
- Domestic Migrants from the South
- Domestic Migrants from the Mid West
- Detailed Job Categories shares
- Detailed Country of Birth shares

Before starting the training model, we follow a standard pre-processing machine learning procedure: we impute missing values with the corresponding median values, and we standardize all our predictors. The random forest is trained on a 75% random sample of the original dataset selected with stratified sampling according to HOLC grades and city population. We set the parameter $m = 41$, which determines the number of variables randomly selected at each split, following the results of an automated tuning procedure employing model-based optimization (MBO) (Probst et al., 2018). The fraction of observations randomly sampled for each tree is grade-specific to counteract the class imbalance of HOLC grades. Hence, the less frequent class (A) has the highest sampling fraction (92%), while the most frequent classes (C and D) have lower fractions (63% and 70%). The results are robust to using a unique sample fraction. In particular the one suggested by the automated MBO procedure of Probst et al. (2018) is 0.89.

A.4 List of Cities

The difference-in-differences results are based on the definition of control and treatment group outlined in Section 4. As a reminder, treated cities are municipalities surveyed by the HOLC with a population between 40,000 and 50,000. The control group includes cities between 30,000 and 40,000 residents that had a distance of least 50km (31mi) from the nearest mapped city.

Treated Cities

- Amarillo, TX
- Dubuque, IA
- Lima, OH
- Portsmouth, OH
- Aurora, IL
- Elmira, NY
- Lorain, OH
- Poughkeepsie, NY
- Chelsea, MA
- Haverhill, MA
- Lynchburg, VA
- Salem, MA
- Chicopee, MA
- Jackson, MS
- Muncie, IN
- S. Petersburg, FL
- Columbus, GA
- Joliet, IL
- Oshkosh, WI
- Stamford, CT
- Council Bluffs, IA
- Lexington, KY
- Phoenix, AZ
- Stockton, CA

- Waterloo, IA
- Woonsocket, RI
- Waltham, MA
- Ogden, UT
- Jackson, MS
- Pueblo, CO
- Warren, OH
- Everett, MA

Control Cities

- Baton Rouge, LA
- Muskogee, OK
- Danville, IL
- Meriden, CT
- Bellingham, WA
- Norwood, OH
- Elkhart, IN
- Green Bay, WI
- Butte, MT
- Paducah, KY
- Everett, WA
- Easton, PA
- Colorado Springs
- Pensacola, FL
- Hazleton, PA
- Fort Smith, AR
- Quincy, IL
- High Point, NC
- Santa Ana, CA
- Hagerstown, MD
- Rock Island, IL
- Marion, OH
- Richmond, IN
- Joplin, MO
- San Bernardino
- Newark, OK
- Sioux Falls, SD
- La Crosse, WI
- Santa Barbara
- Port Huron, MI
- Tucson, AZ
- Laredo, TX
- Alton, IL
- Raleigh, NC
- Lewiston, ME
- Amsterdam, NY
- Rome, NY
- Watertown, NY
- Mansfield, OH
- Auburn, NY
- Sheboygan, WI
- Meridian, MS
- Bloomington, IL
- Steubenville, OH
- Wilmington, NC
- Moline, IL
- Cumberland, MD
- Kokomo, IN
- Zanesville, OH

Table A1 — Neighborhood Distribution according to HOLC Grades

Spatial Unit	N	Proportions			
		A	B	C	D
HOLC Neighborhoods	8872	11.7%	26.3%	38.1%	23.9%
Hexagons	278066	7.8%	22.0%	42.1%	28.0%

Notes: The sample of HOLC neighborhoods includes all the shapes digitized by Nelson et al. (2021) for 202 cities. We obtain the sample of hexagons overlaying a regular grid of hexagons with an area of 0.025 km^2 and a side of approximately 100 *mt.* over the digitized HOLC shapes. The grade of a hexagon is the one occupying the majority of its area. We keep only hexagons whose area is occupied by a single grade for at least 75%.

Table A2 — Census and CoreLogic coverage of Neighborhoods and Cities.

Decade	Share of Coverage				
	Census Data		Year	CoreLogic Deeds	
	Neighborhood	City		Neighborhood	City
1910	45.6%	98.8%			
1920	56.9%	100.0%			
1930	84.5%	100.0%			
1940	96.6%	100.0%			
1950	5.2%	14.8%			
1960	54.4%	51.9%			
			1965	3.4%	35.21%
1970	83.5%	80.2%	1970	5.6%	38.0%
			1975	8.6%	50.7%
1980	87.4%	93.8%	1980	11.4%	57.7%
			1985	21.1%	76.0%
1990	99.4%	100.0%	1990	35.8%	84.5%
			1995	52.7%	88.7%
2000	100.0%	100.0%	2000	64.4%	91.5%
			2005	76.0%	94.4%
2010	100.0%	100.0%	2010	77.9%	94.4%

Notes: The Table reports the percentages of coverage for neighborhoods and cities of interest. The sample includes every hexagon in cities with a 1930 population between 30,000 and 50,000 and at least 20 residents in 1930. CoreLogic deeds are binned in 5-year time periods according to their sale year and month.

Table A3—Short-term Diff-in-Diff Results. No Grade Heterogeneity

	<i>Dependent variables:</i>			
	Homeownership Rates	African American Percentage	Property Values	Rent Prices
DiD	-0.006 (0.007)	0.002 (0.002)	-254* (132)	6.3 (8.1)
\bar{Y}	0.49	0.06	5439	42.6

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade. The regressions are estimated with individual-level observations. The sample includes individuals with valid geocodes in cities with a 1930 population between 30,000 and 50,000, living in hexagons with at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city-year level. See the Notes of Table 7 for additional estimation details.

Significance: * 0.10 ** 0.05 *** 0.01

Table A4—Random Forest Performance, Confusion Matrix, Restricted Test Set

		Data			
		D	C	B	A
Prediction	D	2139	117	12	0
	C	233	4435	219	18
	B	36	214	2647	153
	A	3	18	45	866
Accuracy		90.43%			
Class Sensitivity		88.72	92.70	90.56	83.51
Prevalence		21.61	42.89	26.20	9.30
Detection Prevalence		20.33	43.97	27.34	8.35

Notes: The matrix compares the observed and the predicted grades for a test set of observations excluded from the training procedure. The test set is a 25% random subsample of the original dataset selected with stratified sampling according to city population and HOLC grade. In this case, the test set is restricted to include only hexagons in cities with a population below 50,000. For other details, see the notes of Table 3, which contains the predicted grades for the complete test set.

Table A5 — Logit Performance, Confusion Matrix

		Data			
		D	C	B	A
Prediction	D	8732	2101	118	12
	C	5076	15338	4215	379
	B	158	2801	6225	1998
	A	1	27	423	942
Accuracy		64.35%			
Class Sensitivity		62.52	75.68	56.69	28.20
Prevalence		28.77	41.75	22.62	6.86
Detection Prevalence		22.58	51.51	23.03	2.86

Notes: The matrix compares the observed and the predicted grades for a test set of observations excluded from the training procedure of a logit model. The test set is a 25% random subsample of the original dataset selected with stratified sampling according to city population and HOLC grade. The logit model is estimated with the same estimation steps of the random forest. For other details, see the notes of Table 3, which contains the predicted grades for a random forest algorithm.

Table A6 — 1930 Descriptive Statistics According to Predicted Grades

	Predicted Grade			
	C		D	
	Correct	Wrong	Correct	Wrong
<i>Panel A: 1930 Levels</i>				
Black	0.02 (0.08)	0.02 (0.08)	0.22 (0.34)	0.12 (0.26)
Home Owner	0.53 (0.22)	0.49 (0.24)	0.39 (0.23)	0.43 (0.21)
Income Score	7.13 (0.17)	7.13 (0.16)	6.85 (0.28)	6.92 (0.24)
First Gen Immigrant	0.22 (0.17)	0.26 (0.18)	0.23 (0.23)	0.23 (0.21)
<i>Panel B: 1930-1920 Trends</i>				
Black	-0.005 (0.07)	-0.001 (0.07)	0.03 (0.16)	0.01 (0.19)
Home Owner	-0.01 (0.24)	0.03 (0.23)	-0.02 (0.25)	-0.03 (0.25)
Income Score	0.01 (0.18)	0.03 (0.19)	-0.02 (0.22)	-0.04 (0.25)
First Gen Immigrant	-0.02 (0.14)	-0.03 (0.15)	-0.05 (0.17)	-0.03 (0.14)

Notes: The Table reports averages of 1930 Census variables according to different classifications. The first two columns compare means between hexagons classified as *C* by our classification model. The first columns reports averages for hexagons whose observed grade is *C*, while the second refers to neighborhoods with a HOLC grade other than *C*. The third and fourth columns do the same for grade *D*. The level of observation is a neighborhood (hexagon). The sample includes all the hexagons intersecting a HOLC neighborhood digitized by Nelson et al. (2021) in 202 maps. Standard deviations are reported in parentheses.

Table A7 — Testing Differences in 1920-1910 Trends by Treatment Status

	A	B	C	D
Black	0.00002 (0.004)	0.008 (0.006)	0.007 (0.005)	0.011 (0.014)
Home Owner	0.026 (0.036)	0.009 (0.023)	0.015 (0.010)	0.037* (0.020)
Income Score	0.017 (0.046)	-0.008 (0.014)	-0.003 (0.007)	-0.002 (0.013)
Education Score	2.501 (2.569)	-0.548 (0.528)	0.013 (0.226)	0.235 (0.318)
First Gen Immigrant	-0.025 (0.023)	-0.010 (0.008)	0.013 (0.008)	-0.014 (0.011)
Number of Children	-0.015 (0.028)	0.009 (0.012)	0.009 (0.008)	0.008 (0.014)

Notes: The Table reports the coefficients from a set of regressions where the dependent variable is the 1920-1910 change in the variable reported in the left column, and the independent variable is an indicator for the treatment status. See Appendix Section A.2 for definitions of Census variables in our dataset. The level of observation is a spatial unit (hexagon). The sample includes every hexagon in cities with a 1930 population between 30,000 and 50,000 and at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city level. Significance: * 0.10 ** 0.05 *** 0.01

Table A8 — Short-term Difference-in-Differences Results, by Grade

<i>Dependent variable: African American Percentage</i>					
	(1)	(2)	(3)	(4)	(5)
<i>DiD_A</i>	-0.003	-0.006	-0.004	-0.005	-0.005
$\bar{Y}^A = 0.008$	(0.014)	(0.007)	(0.006)	(0.006)	(0.006)
<i>DiD_B</i>	0.009	0.006*	0.004	0.006*	0.006*
$\bar{Y}^B = 0.01$	(0.012)	(0.003)	(0.003)	(0.004)	(0.004)
<i>DiD_C</i>	0.004	0.002	0.001	0.001	0.001
$\bar{Y}^C = 0.02$	(0.009)	(0.002)	(0.001)	(0.002)	(0.001)
<i>DiD_D</i>	0.012	0.010***	0.014***	0.016***	0.018***
$\bar{Y}^D = 0.19$	(0.096)	(0.003)	(0.004)	(0.004)	(0.004)
City Fixed. Eff.		X		X	X
Spatial Unit Fixed Eff.			X		
Spatial Unit Controls				X	X
Local Area Controls					X

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade. Each row contains the DiD coefficients for a given grade. The regressions are estimated with individual-level observations. The sample includes individuals with valid geocodes in cities with a 1930 population between 30,000 and 50,000, living in hexagons with at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city-year level. The table structure is analogous to Table 7. See the Notes of Table 7 for additional estimation details. Significance: * 0.10 ** 0.05 *** 0.01

Table A9 — Short-term Difference-in-Differences Results, by Grade

<i>Dependent variable: Property Values</i>					
	(1)	(2)	(3)	(4)	(5)
<i>DiD_A</i> $\bar{Y}^A = 22,907$	946 (1,742)	1,219*** (435)	1,142*** (378)	1,239*** (403)	1,046*** (378)
<i>DiD_B</i> $\bar{Y}^B = 16,828$	-178 (604)	-147 (173)	-34 (171)	-106 (178)	-105 (170)
<i>DiD_C</i> $\bar{Y}^C = 12,983$	-510 (383)	-493*** (144)	-487*** (150)	-468*** (148)	-502*** (148)
<i>DiD_D</i> $\bar{Y}^D = 8,615$	-319 (432)	-284* (156)	-301* (157)	-251 (155)	-302** (153)
City Fixed. Eff.		X		X	X
Spatial Unit Fixed Eff.			X		
Spatial Unit Controls				X	X
Local Area Controls					X

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade. Each row contains the DiD coefficients for a given grade. The regressions are estimated with individual-level observations. The sample includes individuals with valid geocodes in cities with a 1930 population between 30,000 and 50,000, living in hexagons with at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city-year level. The table structure is analogous to Table 7. See the Notes of Table 7 for additional estimation details. Significance: * 0.10 ** 0.05 *** 0.01

Table A10 — Short-term Difference-in-Differences Results, by Grade

<i>Dependent variable: Property Values (logs)</i>					
	(1)	(2)	(3)	(4)	(5)
<i>DiD_A</i> $\bar{Y}^A = 9.82$	0.135 (0.244)	0.168*** (0.027)	0.152*** (0.022)	0.168*** (0.026)	0.156*** (0.025)
<i>DiD_B</i> $\bar{Y}^B = 9.52$	-0.057 (0.123)	-0.048* (0.028)	-0.035 (0.028)	-0.040 (0.028)	-0.038 (0.028)
<i>DiD_C</i> $\bar{Y}^C = 9.20$	-0.052 (0.095)	-0.048* (0.027)	-0.049* (0.028)	-0.043 (0.028)	-0.049* (0.028)
<i>DiD_D</i> $\bar{Y}^D = 8.66$	-0.006 (0.190)	-0.011 (0.045)	-0.018 (0.043)	-0.005 (0.044)	-0.022 (0.042)
City Fixed. Eff.		X		X	X
Spatial Unit Fixed Eff.			X		
Spatial Unit Controls				X	X
Local Area Controls					X

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade. Each row contains the DiD coefficients for a given grade. The regressions are estimated with individual-level observations. The sample includes individuals with valid geocodes in cities with a 1930 population between 30,000 and 50,000, living in hexagons with at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city-year level. The table structure is analogous to Table 7. See the Notes of Table 7 for additional estimation details. Significance: * 0.10 ** 0.05 *** 0.01

Table A11 — Short-term Difference-in-Differences Results. Neighborhood Level

<i>Dependent variable: Home-Ownership Rates</i>				
	(1)	(2)	(3)	(4)
DiD_A $\bar{Y}^A = 0.71$	0.034 (0.128)	0.041** (0.017)	0.037** (0.017)	0.037** (0.016)
DiD_B $\bar{Y}^B = 0.65$	0.0002 (0.023)	-0.0001 (0.009)	-0.008 (0.009)	-0.003 (0.009)
DiD_C $\bar{Y}^C = 0.53$	-0.014 (0.029)	-0.015 (0.010)	-0.019** (0.010)	-0.014 (0.009)
DiD_D $\bar{Y}^D = 0.43$	-0.023 (0.046)	-0.019** (0.009)	-0.027*** (0.010)	-0.020** (0.009)
City Fixed. Eff.		X	X	X
Spatial Unit Controls			X	X
Local Area Controls				X

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade. Each row contains the DiD coefficients for a given grade. The regressions are estimated with neighborhood-level observations. The sample includes neighborhoods with at least 20 residents in 1930 in cities with a population between 30,000 and 50,000. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table A12 — Short-term Difference-in-Differences Results. Neighborhood Level

<i>Dependent variable: African American Percentage</i>				
	(1)	(2)	(3)	(4)
DiD_A $\bar{Y}^A = 0.009$	0.005 (0.016)	0.003 (0.006)	0.004 (0.005)	0.004 (0.005)
DiD_B $\bar{Y}^B = 0.01$	0.002 (0.010)	0.001 (0.003)	0.002 (0.003)	0.002 (0.003)
DiD_C $\bar{Y}^C = 0.02$	0.002 (0.009)	0.001 (0.001)	0.0003 (0.001)	0.0005 (0.001)
DiD_D $\bar{Y}^D = 0.22$	0.022 (0.090)	0.015** (0.006)	0.018*** (0.006)	0.018*** (0.006)
City Fixed. Eff.		X	X	X
Spatial Unit Controls			X	X
Local Area Controls				X

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade. Each row contains the DiD coefficients for a given grade. The regressions are estimated with neighborhood-level observations. See the Notes of Appendix Table A11 for additional details. Significance: * 0.10 ** 0.05 *** 0.01

Table A13 — Short-term Difference-in-Differences Results. Neighborhood Level

	<i>Dependent variable: Property Values</i>			
	(1)	(2)	(3)	(4)
<i>DiD_A</i>	629	831	849	747
$\bar{Y}^A = 25,162$	(1,931)	(378)	(354)	(350)
<i>DiD_B</i>	-124	-144	-149	-163
$\bar{Y}^B = 17,371$	(735)	(188)	(189)	(173)
<i>DiD_C</i>	-394	-397***	-348***	-369***
$\bar{Y}^C = 12,727$	(434)	(126)	(129)	(131)
<i>DiD_D</i>	-548**	-523***	-462***	-510***
$\bar{Y}^D = 7,923$	(374)	(143)	(144)	(146)
City Fixed. Eff.		X	X	X
Spatial Unit Controls			X	X
Local Area Controls				X

Notes: The Table reports difference-in-differences coefficients obtained estimating equation 1 by grade. Each row contains the DiD coefficients for a given grade. The regressions are estimated with neighborhood-level observations. See the Notes of Appendix Table A11 for additional details. Significance: * 0.10 ** 0.05 *** 0.01

Table A14 — Short-term Difference-in-Differences Results. Neighborhood Level

	<i>Dependent variable: Rent Prices</i>			
	(1)	(2)	(3)	(4)
<i>DiD_A</i>	60.0	68.8***	66.2***	64.7***
$\bar{Y}^A = 562.20$	(37.4)	(25.2)	(25.0)	(24.8)
<i>DiD_B</i>	18.3	17.6	18.3	18.2
$\bar{Y}^B = 342.76$	(19.4)	(13.3)	(12.6)	(12.5)
<i>DiD_C</i>	22.6	22.3**	21.4**	21.6**
$\bar{Y}^C = 241.72$	(15.5)	(10.4)	(10.3)	(10.3)
<i>DiD_D</i>	6.9	7.5	8.5	7.6
$\bar{Y}^D = 172.63$	(11.5)	(7.5)	(7.5)	(7.5)
City Fixed. Eff.		X	X	X
Spatial Unit Controls			X	X
Local Area Controls				X

Notes: The Table reports difference-in-differences coefficients obtained estimating equation 1 by grade. Each row contains the DiD coefficients for a given grade. The regressions are estimated with neighborhood-level observations. See the Notes of Appendix Table A11 for additional details. Significance: * 0.10 ** 0.05 *** 0.01

Table A15 — Short-term Diff-in-Diff by Grade. City Level S.E. Clustering

<i>Dependent Variables</i>				
	Homeownership Rates	African American Percentage	Property Values	Rent Prices
<i>DiD_A</i>	0.045** (0.021)	-0.005 (0.009)	1,046** (538)	39.3*** (16.3)
<i>DiD_B</i>	-0.002 (0.013)	0.006* (0.005)	-106 (248)	8.8 (17.0)
<i>DiD_C</i>	-0.017 (0.011)	0.001 (0.002)	-502** (212)	5.9 (13.3)
<i>DiD_D</i>	-0.024* (0.013)	0.018*** (0.006)	-302 ** (218)	3.1 (11.4)

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade for four different outcomes. Each row contains the DiD coefficients for a given grade. The Table shows the DiD coefficients resulting from a DiD framework with a city fixed effect and geographic and demographic controls at the hexagon level. The regression specification is analogous to the one in column (4) of Table 7. The list of controls includes geographic coordinates, a scaled measure of distance from the city center, spatial unit's population density, imputed income score and family size. The regressions are estimated with individual-level observations. The sample includes individuals with valid geocodes in cities with a 1930 population between 30,000 and 50,000, living in hexagons with at least 20 residents in 1930. See Appendix Section A.4 for a list of cities. Standard errors, in parentheses, are clustered at the city level. Significance: * 0.10 ** 0.05 *** 0.01

Table A16 — Short-term Diff-in-Diff. Grouping with Observed Grades

<i>Dependent Variables</i>			
	Homeownership Rates	African American Percentage	Property Values
<i>DiD_A</i>	0.034** (0.015)	-0.004 (0.004)	1,253*** (360)
<i>DiD_B</i>	0.004 (0.008)	0.006* (0.003)	-171 (187)
<i>DiD_C</i>	-0.010 (0.007)	0.002 (0.002)	-407*** (144)
<i>DiD_D</i>	-0.018* (0.008)	0.018*** (0.004)	-405 ** (167)

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade for three different outcomes. Each row contains the DiD coefficients for a given grade. Differently from the other results, we group observations in treated cities according to their observed HOLC grade, rather than the predicted one. See the Notes of Table A15 for estimation details. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table A17 — Short-term Diff-in-Diff. Extended Treatment Group

	<i>Dependent Variables</i>		
	Homeownership Rates	African American Percentage	Property Values
<i>DiD_A</i>	0.027 (0.019)	0.008** (0.004)	1,657** (835)
<i>DiD_B</i>	0.001 (0.008)	0.001 (0.002)	-766*** (255)
<i>DiD_C</i>	-0.017** (0.007)	0.0002 (0.001)	-578*** (130)
<i>DiD_D</i>	-0.024*** (0.008)	0.019*** (0.006)	-855 *** (192)

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) at the neighborhood level, by grade, for three different outcomes. Each row contains the DiD coefficients for a given grade. Differently from the other results, we extend the treatment group to include cities between 40,000 and 60,000 residents in the 1930 census. See the Notes of Table A15 for estimation details. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table A18 — Short-term Diff-in-Diff. Placebo Outcomes

<i>Dependent Variables</i>			
	Female Percentage	Number of Children	Unemployment Rate, Men
DiD_A	0.004 (0.002)	0.019 (0.013)	-0.045*** (0.008)
DiD_B	0.005*** (0.002)	0.004 (0.002)	0.011** (0.005)
DiD_C	-0.001 (0.001)	-0.006 (0.005)	-0.003 (0.007)
DiD_D	0.002 (0.002)	0.022*** (0.006)	0.012 (0.013)
\bar{Y}^A	0.523	0.824	0.070
\bar{Y}^B	0.527	0.744	0.059
\bar{Y}^C	0.509	0.778	0.099
\bar{Y}^D	0.495	0.821	0.136

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) by grade for three different outcomes. Each row contains the DiD coefficients for a given grade. See the Notes of Table A15 for estimation details. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table A19 — Short-term Diff-in-Diff. Placebo Population Threshold

	<i>Dependent Variables</i>		
	Homeownership Rates	African American Percentage	Property Values
<i>DiD_A</i>	0.016 (0.026)	0.004 (0.004)	-1,999 (1,498)
<i>DiD_B</i>	-0.015 (0.014)	-0.003* (0.002)	717 (492)
<i>DiD_C</i>	-0.008 (0.010)	0.001 (0.002)	53 (301)
<i>DiD_D</i>	0.004 (0.008)	0.009 (0.006)	970*** (324)

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) at the neighborhood level, by grade, for three different outcomes. Each row contains the DiD coefficients for a given grade. Differently from the other results, we set a placebo population threshold at 60,000. Accordingly, we define neighborhoods in cities between 60,000 and 70,000 residents as treated, while areas in cities between 50,000 and 60,000 are included in the control group. See the Notes of Table A15 for estimation details. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table A20 — Short-term Diff-in-Diff. Neighborhood Grade Index.

	<i>Dependent Variables</i>		
	Homeownership Rates	African American Percentage	Property Values
<i>DiD_A</i>	0.042*** (0.015)	-0.005 (0.006)	1,223*** (398)
<i>DiD_B</i>	-0.002 (0.009)	0.006* (0.004)	-106 (177)
<i>DiD_C</i>	-0.017** (0.007)	0.001 (0.002)	-472*** (147)
<i>DiD_D</i>	-0.022** (0.009)	0.015*** (0.004)	-274* (155)

Notes: The Table reports difference-in-differences coefficients obtained estimating equation (1) at the neighborhood level, by grade, for three different outcomes. Each row contains the DiD coefficients for a given grade. Differently from the other results, we include the grade composition of the surrounding neighborhoods as control. The surrounding neighborhoods are defined with a 1000mt radius (0.63 miles). See the Notes of Table A15 for estimation details. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Table A21 — Long-Term Difference-in-Differences Results, by Grade. CoreLogic

<i>Dependent variable: Property Values (logs)</i>				
Grade				
	A	B	C	D
DiD_{65}	. (.)	0.05 (0.10)	-0.52*** (0.09)	-0.95*** (0.13)
DiD_{70}	-0.24 (0.17)	0.025 (0.121)	-0.68*** (0.12)	-1.05*** (0.14)
DiD_{75}	-0.15 (0.17)	-0.010 (0.113)	-0.53*** (0.10)	-0.94*** (0.15)
DiD_{80}	-0.17 (0.17)	-0.11 (0.12)	-0.77** (0.19)	-0.70*** (0.22)
DiD_{85}	0.18 (0.23)	0.19 (0.18)	-0.09 (0.22)	-0.51** (0.25)
DiD_{90}	0.29 (0.19)	0.23* (0.18)	-0.06 (0.18)	-0.40 (0.249)
DiD_{95}	-0.02 (0.21)	0.02 (0.16)	-0.28* (0.15)	-0.55*** (0.19)
DiD_{00}	0.07 (0.23)	0.06 (0.16)	-0.14 (0.18)	-0.39** (0.19)
DiD_{05}	0.21 (0.25)	-0.06 (0.19)	-0.13 (0.19)	-0.35 (0.21)

Notes: The Table reports Difference-in-Differences coefficients obtained estimating equation (2) by grade. Each row contains the DiD coefficients for a given grade in the corresponding year. The regressions are estimated with neighborhood-level observations. The sample includes neighborhoods with at least 20 residents in 1930 in cities with a population between 30,000 and 50,000. See Appendix Section A.4 for a list of cities. The data source for post-1940 outcomes is CoreLogic, see Section 3.3 for details. The outcome variable is adjusted with CPI to 1980 dollars. Standard errors, in parentheses, are clustered at the city-year level. Significance: * 0.10 ** 0.05 *** 0.01

Observable Pre-Trends, Grades A and B

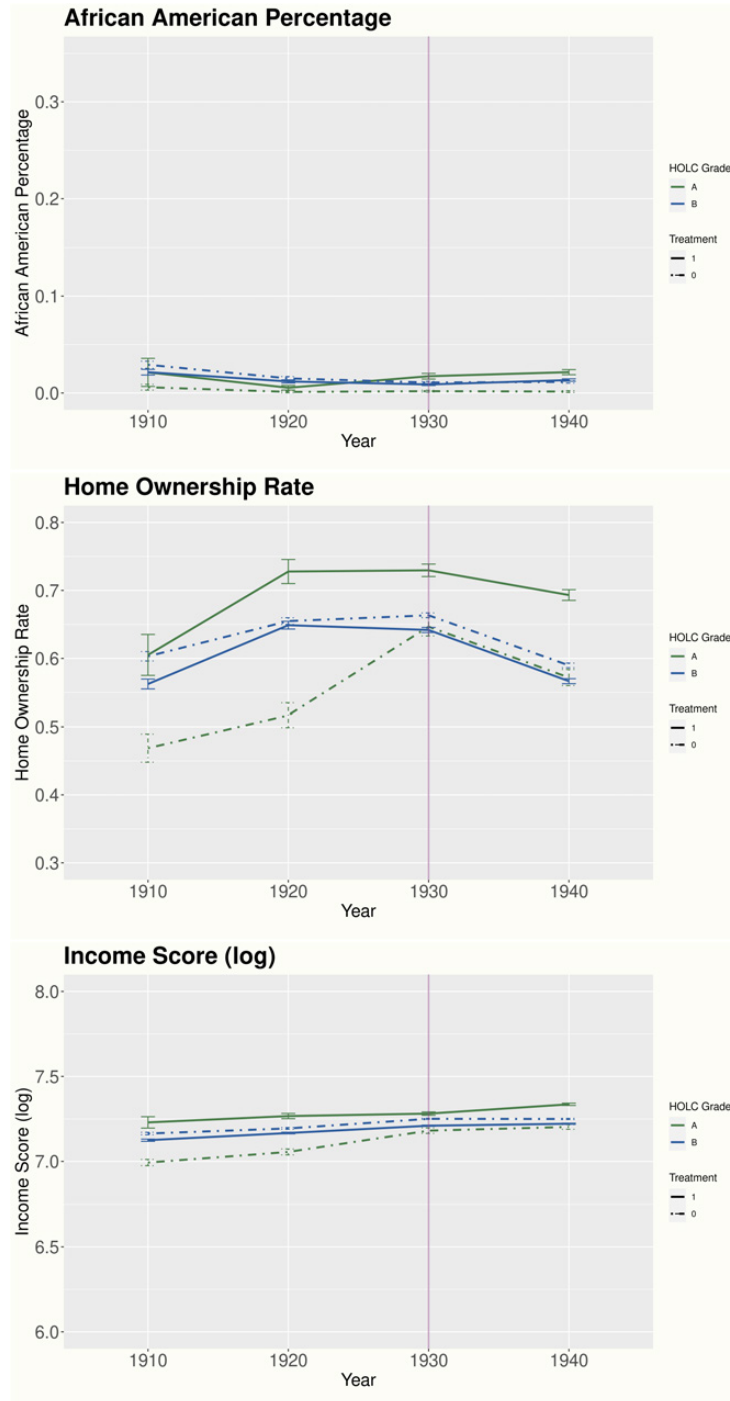


Figure A1 — The Figure shows pre-trends for selected variables for A and B grades. The point estimates are averages of hexagon-level observations. The bars show the respective standard errors of each mean. The sample includes neighborhoods in cities with a 1930 population between 30,000 and 50,000, with at least 20 residents in 1930. The vertical line highlights 1930, the last pre-treatment decade. See Appendix Section A.2 for definitions of Census variables in our dataset.

Observable Pre-Trends, Additional Variables

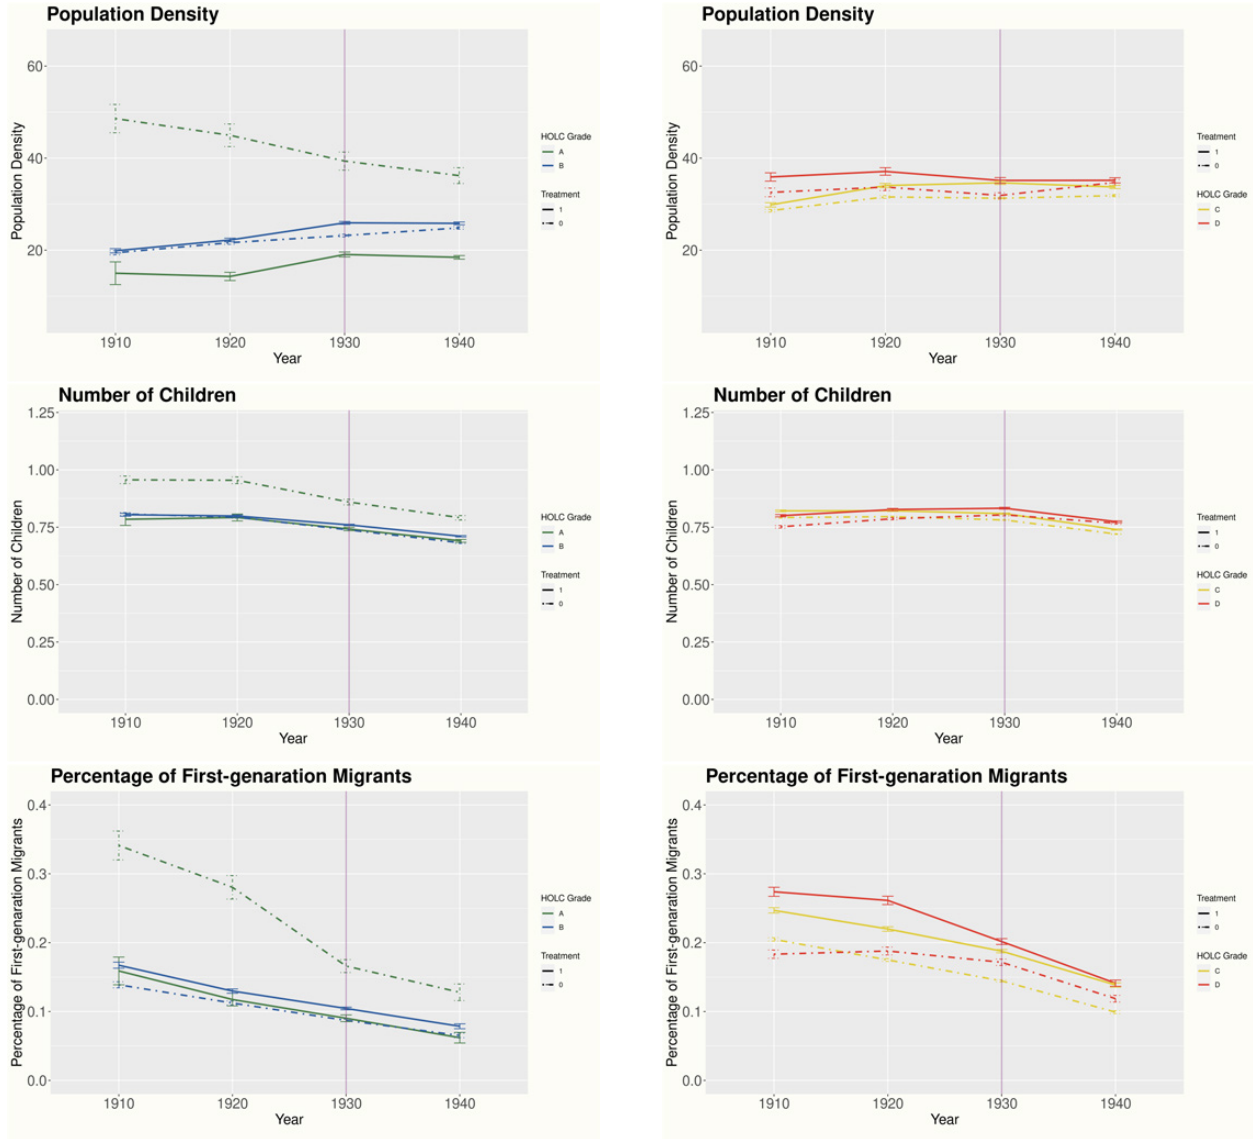


Figure A2— The Figure shows pre-trends for selected variables according to their predicted grade. The point estimates are averages of hexagon-level observations. The bars show the respective standard errors of each mean. The sample includes hexagons in cities with a 1930 population between 30,000 and 50,000, with at least 20 residents in 1930. The vertical line highlights 1930, the last pre-treatment decade. See Appendix Section A.2 for definitions of Census variables in our dataset.

Accuracy Levels according to Dataset Definition




Figure A3 — The Figure shows the accuracy level we obtain when we train the random forest classification algorithms with different datasets according to the size of cities we include. The purple line plots the Accuracy obtained with a test set defined as a 25% random subsample of the original dataset, selected with stratified sampling according to city population and HOLC grade. The brown line shows Accuracy levels when we restrict the test set to cities with at most 50,000 residents. The level of observation is a neighborhood (hexagon). See Section 3.1 for details about the hexagon definition. The complete dataset includes every hexagon in a mapped city containing at least 20 residents in 1930. Overall Accuracy is the percentage of hexagons whose predicted grades correspond to observed ones. A predicted grade is the class predicted by the trained random forest algorithm. See Section 4.1 and Appendix Section A.3 for details about the Random Forest training procedure.

Long-Term Difference-in-Differences Results, A and B Grades. Census Data




Figure A4 — The Figure shows the estimated coefficients for regression (2) and their 95% confidence intervals for grades A and B. The coefficients and standard errors are analogous to the ones reported in Table 9, but focus on different grades. See the Notes of Table 9 for estimation details.

Long-Term Difference-in-Differences Results, A and B Grades. CoreLogic Data




Figure A5— The Figure shows the estimated coefficients for regression (2) and their 95% confidence intervals for property values. The Figure includes the results for grades C and D. The coefficients and standard errors are the ones reported in Table 10. See the Notes of Table 10 for estimation details.

Long-Term Difference-in-Differences Results, by Grade. CoreLogic Data




Figure A6— The Figure shows the estimated coefficients for regression (2) and their 95% confidence intervals for property values. The coefficients and standard errors are analogous to the ones reported in Table 10, but the outcome is a log-transformation. See the Notes of Table A21 for estimation details.