The impact of 1930s maps of mortgage risk on modern home ownership

Randy Cragun[a,b] and Lauren Barnett[b]

[a]Corresponding author: rscragun@gmail.com or rc@rcragun.net.

[b]Birmingham-Southern College, 900 Arkadelphia Road, Birmingham, AL, 35254, USA

August 30, 2024

Abstract

This research estimates long-term impacts of "redlining" maps on rates of home ownership. In the 1930s, the Home Owners' Loan Corporation drew maps that split zones in American cities into 4 categories of mortgage risk. We exploit the grade discontinuities at the borders between zones to estimate effects of the grade assignments. We match Census blocks to HOLC grades and then build small neighborhoods of blocks that are close together and lie along both sides of HOLC borders. We estimate impacts of HOLC grades by comparing outcomes in 2010 between higher- and lower-graded blocks within neighborhoods. Estimates suggest substantial long-term impacts of receiving a lower HOLC grade, with people in lower-graded blocks being between 2 and 13 percentage points less likely to live in an owner-occupied dwelling than their neighbors a couple blocks away in a higher-graded zone.

Keywords: Housing, HOLC, Redlining, Homeownership, Spatial discontinuity, Mortgage risk, Racism, Discrimination, Urban economics

JEL Classification: R31, R38, G21, N92, J15

1 Introduction

This research estimates long-term impacts of "redlining" maps on rates of home ownership. The Home Owners' Loan Corporation (HOLC) drew these maps in the 1930s. The maps graded mortgage risk in zones of American cities from A ("Best") to D ("Hazardous"). We match Census blocks to HOLC grades and then build small neighborhoods of blocks that are close together and lie along both sides of HOLC zone borders. We estimate impacts of HOLC grades with variation within neighborhoods.

HOLC grades likely measured many things including access to services, land quality, population density, age of buildings, physical beauty of the space, racial and cultural prejudices against residents, racial homogeneity, and amore We attempt to separate effects of ratings from effects of these other variables using a spatial discontinuity design. HOLC grades vary discontinuously across HOLC zone boundaries (e.g., the grade jumps from C to B suddenly). If unobserved factors that impact outcomes today varied continuously across the HOLC zone boundaries, then effects of the grade can theoretically be estimated consistently by comparing outcomes arbitrarily close to the boundary. In practice, how close we can get to the boundaries is limited by the data.

A contribution of this research is that we use data at the Census block level, whereas past research uses tract-level data. We also implement a novel method of grouping areas into neighborhoods.

Racial discrimination is an important part of the redlining literature and of historical lending practices. However, we do not attempt to discriminate effects due to racial prejudice from effects due to actual mortgage risk. Nor do we estimate effects of being denied a mortgage because of race. However, zones graded as higher risk were more likely than others to be majority Black or to have high concentrations of

minority residents. Thus, if lower HOLC grades in the 1930s decrease rates of home ownership in today, that is important for racial justice. Furthermore, to the extent that HOLC grade differences were due to racial prejudice, effect of those grades today can also be attributed to that prejudice.

It has been almost a century since the HOLC drew its infamous maps. Why would we see effects today? The most obvious way for effects to propagate from the 1930s to today is through the accumulation (or non-accumulation) of family and household wealth in the form owned property. We should expect to find effects today if the people who live in affected neighborhoods today are descended from families that lived in those neighborhoods in the past. However, that is not the only possible mechanism. Household wealth spills over into the surrounding community (for example, through peer effects in school). Households that move into the neighborhood benefit from these spillovers, increasing the probability that they, too, can purchase homes.[1]

2 Background on the HOLC and redlining

The HOLC was a federal agency established in 1933 to expand home-buying opportunities through mortgage refinancing. Although the organization existed to improve opportunities for home ownership, there is some evidence that their practices systematically disadvantaged minorities. The HOLC drew maps that graded zones in US cities by mortgage risk. Lending practices based on these maps are often referred to as "redlining" because the lowest-rated zones were marked in red. These zones received less favorable treatment from the HOLC, the Federal Housing Administration (FHA), and private lenders. The term also gets used more generally for any practice

[1] An interesting study might follow over time the households and families of people who lived near the borders between HOLC zones when the maps were in use. These households experienced a shock to their wealth in the 1930s that might persist as differences in the wealth of their descendants.

of refusing services to persons due to the characteristics of the neighborhoods in which they reside.

The evidence for racial discrimination in HOLC and FHA practices is substantial. The 1938 FHA underwriting manual, for example, instructed FHA agents to consider adverse influences in their appraisal of a property; some of the cited adverse influences included surrounding neighborhoods that had "incompatible racial or social groups" (Federal Housing Administration, 1938, section 937). HOLC auditors' notes justifying their ratings also frequently cite the presence of undesirable racial and ethnic groups or immigrants. The notes on a neighborhood next to Birmingham-Southern College lauded the "proximity to schools, churches, parks, and community centers" before warning about "proximity to negro property on eastern and northern sides of area" (Nelson and et al, 2023).

In later decades, neighborhoods that were redlined experienced increased rates of racial segregation, crime, adverse health risks, rent, etc. and decreased home values and home ownership rates (Aaronson et al., 2021a,b; Anders, 2023; Appel and Nickerson, 2016; Asher, 2021). These disparities are often attributed to the HOLC, both within scholarly discourse and public understandings (Fishback et al., 2022). Existing economic research suggests the creation of the HOLC maps increased crime rates, racial segregation, neighborhood risk factors, and educational deficiencies (Aaronson et al., 2021b; Anders, 2023; Appel and Nickerson, 2016). Further, researchers argue that due to a lack of credit access among communities that received the lowest grades from the HOLC, there are significant differences in property values, home ownership rates, and community disinvestment (Aaronson et al., 2021b; Appel and Nickerson, 2016).

However, recent research questions whether lending companies actually referenced the HOLC maps to make decisions about loaning to a particular individual or if the

maps merely reflected pre-existing trends in lending patterns (Fishback et al., 2022; Xu, 2022). Fishback et al. (2022) and Xu (2022), most notably, argue that the HOLC maps had little causal effect on long-term outcomes; instead, they attribute long-run disparities to practices by the FHA. To reach this conclusion, both Fishback et al. and Xu use unique data collection strategies to conduct their research, as the FHA destroyed almost all of their original records that described the locations in which they had insured mortgage loans (Fishback et al., 2022; Xu, 2022). Xu, for example, is a case study of Chicago, Illinois, as "this is only known available single-family mortgage map" created by the FHA (Xu, 2022, p. 250).

3 Data and methods

Our methodology exploits the spatial discontinuity in treatment with HOLC grades at the boundaries between HOLC zones.

We first assign grades to blocks and then build small "neighborhoods" of blocks that contain blocks of different grades that are close together. We compare outcomes between areas with different HOLC grades within neighborhoods and aggregate those comparisons across neighborhoods.

Getting unbiased estimates of causal effects of HOLC grades requires that the placement of HOLC boundaries within neighborhoods is as good as random. In practical terms, though, the bias will be negligible if pre-treatment factors vary continuously across the boundaries within neighborhoods and our neighborhoods are sufficiently small that these differences are negligible.

3.1 Outcome measure: Home ownership from Census data

We use data at the block level from the 2010 decennial Census—Summary File 1 from 2010 (U.S. Census Bureau, 2010). We drop from the sample any block with 0 reported residents.

Our outcome measure is the proportion of residents in a Census block who live in an owner-occupied dwelling. To measure home ownership, we divide the number of residents in owner-occupied dwellings in a block by the total number of residents in the block.

3.2 Treatment measure: HOLC maps of mortgage risk

We combine the block-level Census data with geocoded versions of HOLC maps from the University of Richmond's "Mapping Inequality" database (Nelson and et al, 2023). Mapping Inequality includes 7,499 HOLC zones in 152 US cities across 29 states. The HOLC gave each zone a grade: A "Best," B "Still Desirable," C "Definitely Declining," or D "Hazardous." We compare areas with A grades to areas with B grades, areas with B grades to areas with C grades, and areas with C grades to areas with D grades.

Approximately 10% of zones have some overlap with another zone. These areas of overlap are small relative to the areas of the zones. If the overlapping zones have the same grade, we combine them into a single zone. If they have different grades, we remove the area of overlap from both zones.

3.3 Assigning grades to blocks

For a block to be included in our sample, we need to assign an HOLC grade to it. If at least 80% of a block's area is contained within HOLC zones of a particular grade,

the block gets that grade.[2] Otherwise, the block is dropped from the sample.

3.4 Neighborhoods

Neighborhoods are specific to a pair of consecutive grades. In other words, we form neighborhoods and conduct the analysis separately for "A" versus "B", "B" versus "C", and "C" versus "D" grades. The process for grouping blocks into neighborhoods is the same for each grade pair, so we present only the process for "C"- and "D"-graded blocks for brevity.

3.4.1 Distance between blocks

To form C-D neighborhoods, we start by finding the closest C block to each D block and the closest D block to each C block. For each possible pair of one C block and one D block within a given state, we calculate the maximum Hausdorff distance between them:

- For each point on the boundary of the C block, we find the distance to the closest point in the D block.
- For each point on the boundary of the D block, we find the distance to the closest point in the C block.
- The maximum Hausdorff distance is the greatest of these distances.

In other words, the distance is small when every point in both blocks is close to some point in the other block. If the distance is greater than some cutoff, we drop that pairing. If a block has no remaining pairings, we drop it from the sample.

[2]This 80% overlap does not have to come from a single HOLC zone. For example, the HOLC could have drawn two C-graded zones close together, and a block could overlap with both, and the block would be included. However, census blocks are quite small relative to HOLC zones, so this is rare (we have not identified any cases).

The distance cutoff is arbitrary, so we report results for cutoffs from 100 m to 600 m in 50 m intervals. Smaller cutoffs have the benefit of reducing bias due to factors that vary with distance to the boundary but reduce the number of data. Using smaller distance cutoffs also tends to drop larger blocks. This is good because non-residential blocks tend to be larger but bad because it makes the sample more urban. Figure 1 shows the distribution of block sizes in the B-C sample as a function of the distance cutoff. That tradeoff has questionable value at higher cutoffs and at very low cutoffs becomes a big problem. The A-B and C-D samples are similar. With a distance cutoff of 600 m, the median block size in our B-C sample is 14862.5 m^2. Suppose such blocks were squares in a grid. A rectangle of two blocks would have a maximum Hausdorff distance of 172.4 m. Squares in a grid are the best case; other shapes and layouts increases that distance. If we set the distance cutoff to 150 m, these median-sized blocks get dropped from the sample.

Using the Hausdorff distance is also somewhat arbitrary. For example, we could use the minimum distance between any two points in the two blocks. Using the Hausdorff distance has the benefit of guaranteeing that everyone who lives in one block lives close to everyone in the other block.

3.4.2 Forming neighborhoods

Neighborhoods are groupings of blocks such that every C or D block is in exactly 1 neighborhood, every C block has its closest D block in the same neighborhood, every D block has its closest C block in the same neighborhood, and the neighborhood cannot be split into two neighborhoods that meet the above conditions (i.e., neighborhoods are as small as possible).

We use the following algorithm to form neighborhoods.

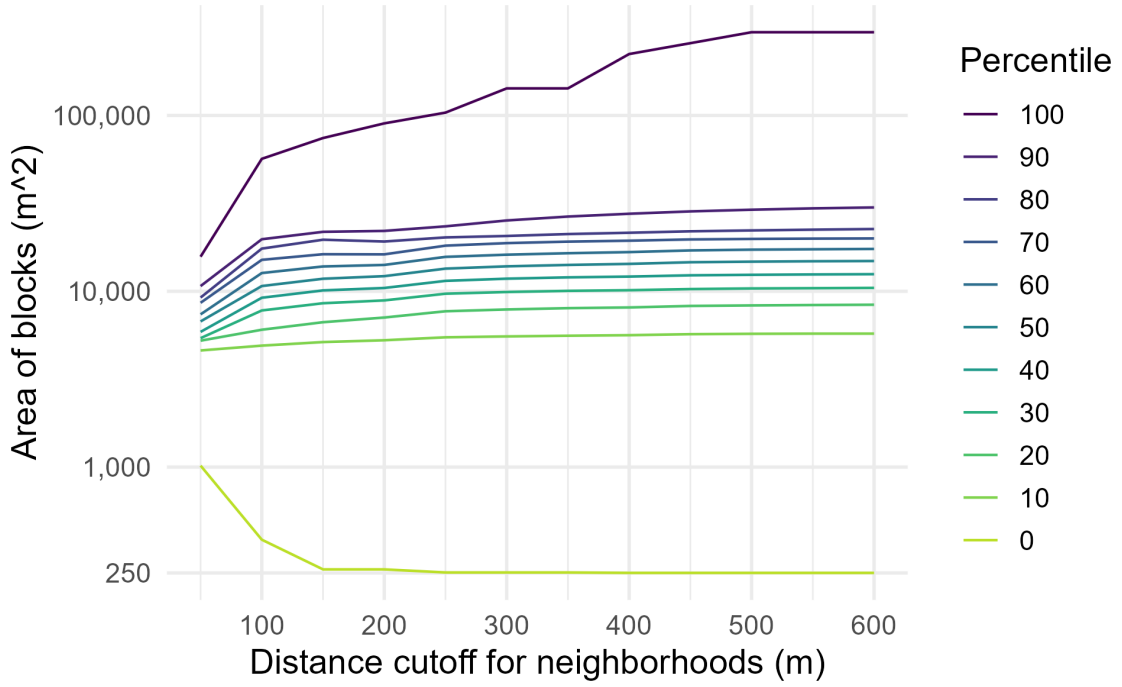

Figure 1: Distribution of block sizes in the B-C sample as a function of the distance cutoff for neighborhoods

- Randomly select one block ("initial block") and assign it to a neighborhood.
- The initial block becomes the "current block".
- Find the block of the opposite grade that is closest to the current block and assign it to the neighborhood.
- That closest block becomes the "current block" (if there are multiple, they all get processed as the current block in sequence).
- Find the closest block for the new current block and add it to the neighborhood.
- Continue assigning blocks to this neighborhood until the closest block is already assigned to the neighborhood.

- If at any step, the closest block is already assigned to a different neighborhood, merge the neighborhoods.
- Randomly select an unassigned block and do the process again.
- Stop when all blocks are assigned.

3.5 Typical levels of home ownership

Figure 2 shows the proportion of people in each of our samples that lived in an owner-occupied dwelling as attached points. Horizontal dashed lines are proportions within the full set of blocks that were assigned a grade regardless of whether they were close enough to other blocks to make it into a neighborhood sample.

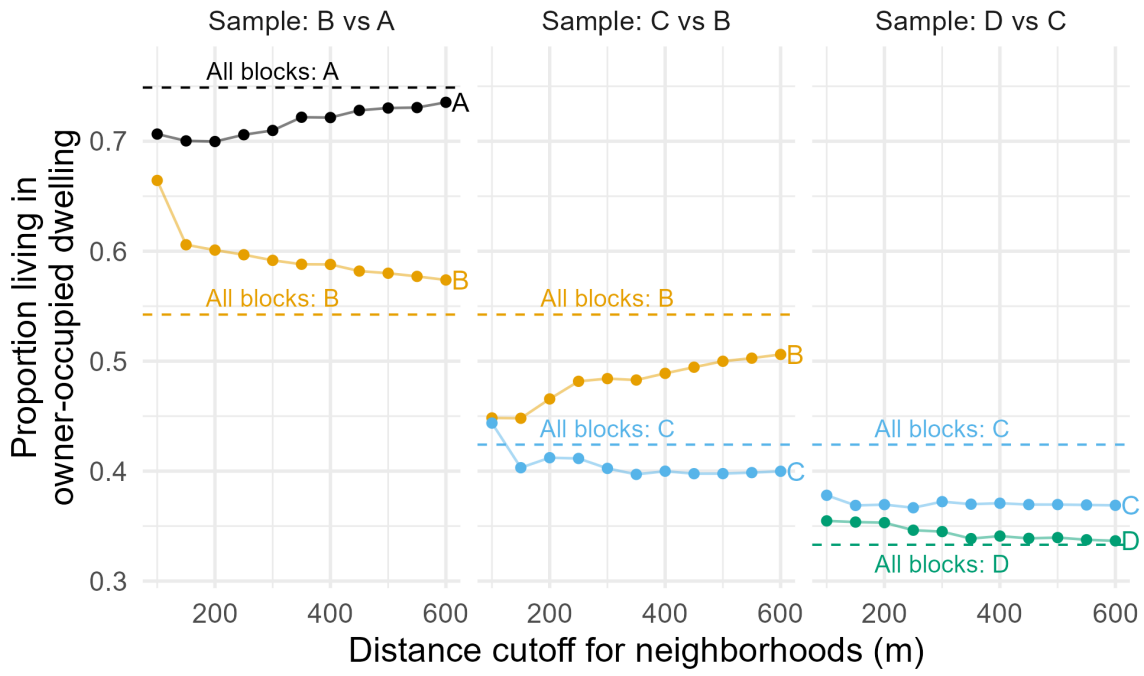

Figure 2: Proportion of residents living in an owner-occupied dwelling by neighborhood distance cutoff. Horizontal dashed lines are proportions within the full set of blocks that were assigned a grade regardless of whether they were close enough to other blocks to make it into a sample.

How representative are our neighborhood samples? In most cases, the rates of home ownership in our samples become more similar to the rates in all graded blocks as we use a larger distance cutoff. All else equal, we would like our samples to be representative of the set of all blocks, but our spatial discontinuity design means we estimate partial effects within the sets of blocks near HOLC boundaries, which may not be representative of all blocks. Differences between graded blocks (mostly urban) and blocks that never got HOLC ratings (mostly rural) may be even more important than the differences shown here between our final samples and the set of all graded blocks

For a given sample, the rates of home ownership tend to fall between the rates for the relevant grades within the full set of blocks. As we might expect, B-graded blocks that are close to A-graded blocks (first panel) have higher rates of home ownership than B-graded blocks do overall. On the other hand, being close to C-graded blocks (second panel) drives down the rates of home ownership among B-graded blocks. Surprisingly, the middle panel shows that people in C-graded blocks near B-graded blocks are actually slightly *less* likely to live in an owned home than people in other C-graded blocks. This result deserves additional study.

Note that the values in Figure 2 are aggregates across all cities in the sample. Our estimates below instead compare home ownership only *within* neighborhoods.

3.6 Estimation procedure

Using the block-level data, we regress the outcome (the proportion of residents living in an owner-occupied dwelling) on an indicator for having the lower grade (e.g. D for a C-D comparison) and fixed effects for neighborhood. Observations are weighted by the number of people in the block because we are interested in the typical experience

of a person. Standard errors are "clustered" by neighborhood (corrected for arbitrary correlation of errors within neighborhoods).

In addition to the clustered standard errors, we tested a randomization inference procedure that estimates the sampling distribution of the estimator under the null hypothesis that HOLC grade has no effect and assuming that HOLC grade was randomly assigned to block within neighborhoods. We do not report the results of these simulations because there is no substantive difference between the p-values calculated from such simulations and those implied by the estimates we report below.

4 Results

Using the block-level data, we regress the rate of home ownership on an indicator for having a the lower grade with a pair of grades and fixed effects for neighborhood. Observations are weighted by the number of residents in the block. Standard errors are "clustered" by neighborhood (corrected for arbitrary correlation of errors within neighborhoods).

The estimates are in Figure 3, which also demonstrates how the estimates differ with the distance cutoff for neighborhoods. Vertical lines are 95% confidence intervals using standard errors clustered at the neighborhood level.

Take the third triangle-shaped point from the left, for example. This point says that when we construct neighborhoods with the distance cutoff of 200 m, people in C-graded blocks are about 5 percentage points less likely to live in an owner-occupied dwelling than people in B-graded blocks in the same neighborhood.

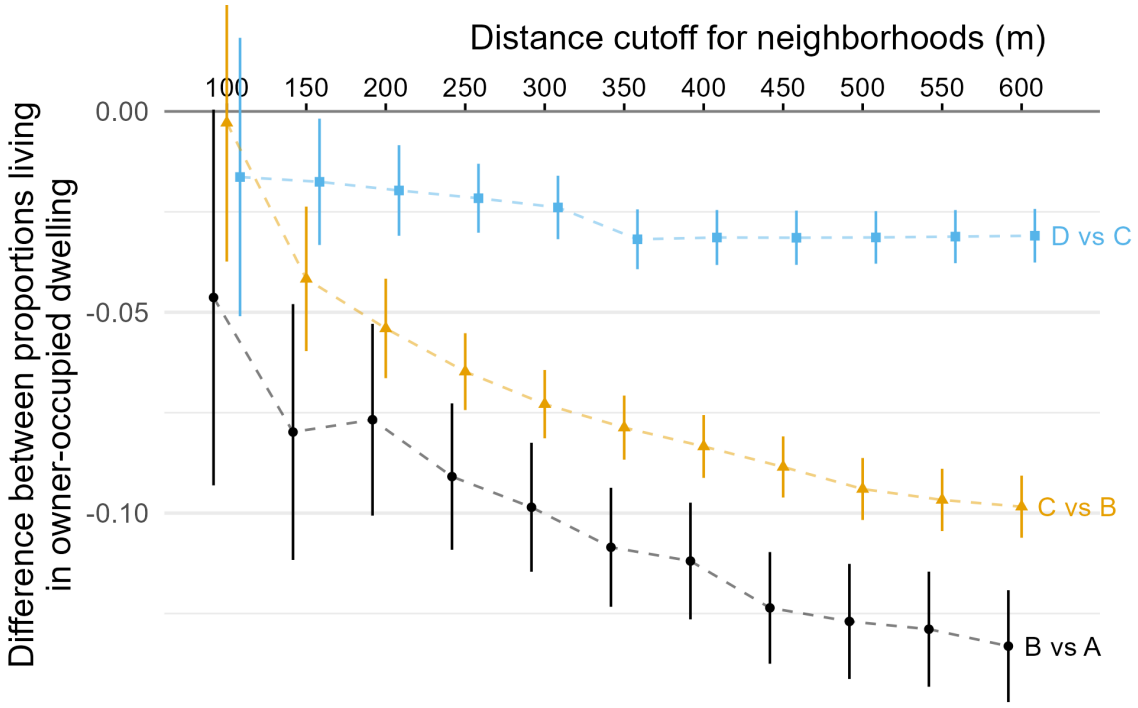

Figure 3: Regression estimates of effects of living in a lower-grade neighborhood on home ownership by neighborhood distance cutoff. Vertical lines are 95% confidence intervals using standard errors clustered at the neighborhood level.

5 Discussion

The findings of this study suggest relationships between HOLC grade and home ownership. All estimates compare the lower-graded area to the higher-graded area, and all are negative as expected. Among people living today near the boundaries between regions given different grades by the HOLC, people on the lower-ranked side are less likely to live in a resident-owned home than their neighbors on the other side of the boundary are.

Like Aaronson et al. (2021a) and Aaronson et al. (2021b), we find larger differences between higher-rated areas (e.g. C versus B) than between lower-rated areas (D versus C).

We get estimates of effect sizes between 2 and 13 percentage points. Such effects are substantial. The estimates in Figure 3 explain about half of the unconditional differences in the full set of all graded blocks in Figure 2.

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.

Aaronson, Daniel, Daniel Hartley, and Bhashkar Mazumder. 2021b. The effects of the 1930s HOLC "redlining" maps. *American Economic Journal: Economic Policy* 13:355–92.

Anders, John. 2023. The long run effects of de jure discrimination in the credit market: How redlining increased crime. *Journal of Public Economics* 222.

Appel, Ian and Jordan Nickerson. 2016. Pockets of poverty: The long-term effects of redlining. Available at SSRN 2852856.

Asher, Lila. 2021. The relationship between historical redlining and Census Bureau Community Resilience Estimates in Columbus, Ohio. *Environment and Planning A: Economy and Space* 53:1859–1861.

Federal Housing Administration. 1938. *Underwriting manual*. Washington, D.C.: U.S. Government Printing Office, revised feb. 1938 ed. FHA Form No. 2049.

Fishback, Price, Jonathan Rose, Ken Snowden, and Thomas Storrs. 2022. New evidence on redlining by federal housing programs in the 1930s. *Journal of Urban Economics* .

Nelson, Robert K. and LaDale Winling et al. 2023. Mapping inequality: Redlining in new deal america. digital scholarship lab. `https://dsl.richmond.edu/panorama/redlining/`.

Xu, Wenfei. 2022. Legacies of institutionalized redlining: a comparison between speculative and implemented mortgage risk maps in Chicago, Illinois. *Housing Policy Debate* 32:249–274.