

The Assessment Gap: Racial Inequalities in Property Taxation

Carlos Avenancio-León, Indiana University

Troup Howard, University of California, Berkeley*

February 2020

Job Market Paper

[Click For Most Recent Version](#)

Abstract

We use panel data covering 118 million homes in the United States, merged with geolocation detail for 75,000 taxing entities, to document a nationwide “assessment gap” which leads local governments to place a disproportionate fiscal burden on racial and ethnic minorities. We show that holding jurisdictions and property tax rates fixed, black and Hispanic residents nonetheless face a 10–13% higher tax burden for the same bundle of public services. This assessment gap arises through two channels. First, property assessments are less sensitive to neighborhood attributes than market prices are. This generates racially correlated spatial variation in tax burden within jurisdiction. Second, appeals behavior and appeals outcomes differ by race. This results in higher assessment growth rates for minority residents. We propose an alternate approach for constructing assessments based on small-geography home price indexes, and show that this reduces inequality by at least 55–70%.

*Corresponding author, Haas School of Business (trouphoward@berkeley.edu). We would like to thank Abhay Aneja, Steve Cicala, Hilary Hoynes, Paulo Issler, Maris Jensen, Andrew Kahrl, Pat Kline, Ross Levine, Ulrike Malmendier, Conrad Miller, Enrico Moretti, Adair Morse, Hoai-Luu Nguyen, Christine Parlour, Sarah Resnick, Rob Ross, Emmanuel Saez, Nick Sander, Nancy Wallace, Danny Yagan, and Gabriel Zucman for their helpful comments. Particular thanks are due to David Sraer for his advice and guidance. Financial support is gratefully acknowledged from the Fisher Center for Real Estate and Urban Economics at Berkeley-Haas. All remaining errors are our own.

1 Introduction

In the United States, the residential property tax is an ad valorem tax. The amount levied should be proportional to the value of the home. Authorizing legislation regularly makes explicit that the relevant concept of value is the market price of the property in a fair transaction. Property tax bills, however, are generated by applying the locally determined rate of taxation to an *assessed value*, which is a local official’s projection of market price. Any wedge between market values and assessed values, therefore, generates some deviation from the intended rate of taxation. Equitable property tax administration requires the ratio of assessed value to market value to be the same for all residents within any particular taxing jurisdiction. This paper documents the existence of a widespread and large racial assessment gap: relative to market value, assessed values are significantly higher for minority residents. This assessment gap places a disproportionate fiscal burden on minority residents: within the same tax jurisdiction, black and Hispanic residents bear a 10–13% higher property tax burden than white residents.

We exploit a property-level dataset spanning most properties in the US, along with a comprehensive record of property transactions assembled from administrative data. We form assessment ratios by restricting the sample to homes for which we observe an assessment and a full-consideration, arm’s-length sale within the same year. Using transacted prices ensures we construct the assessment ratio with an accurate measure of a home’s fair market value. Because property taxes are levied by a wide range of local entities, which often only have partial geographic overlap, it is crucial to compare assessment ratios *within the same tax jurisdiction*. This also helps resolve a practical feature of property taxation, which is that assessments are rarely intended to be one-to-one with market value. The assessing entity chooses a scaling factor for assessments, which can range from less than 10% to 100%, and may change from one year to the next. As a result, variation in assessment ratios across tax jurisdictions may reflect either different levels of taxation or different scaling factors. To address this issue, we exploit a set of shapefiles that provide geographic delineation for the universe of local governments and other taxing entities in the U.S. We use these shapefiles to create unique taxing jurisdictions: properties belonging to the same jurisdiction face the same level of intended taxation, the same set of entities providing public services, and the same assessment scaling factor.

Our main empirical exercise compares assessment ratios – the ratio of a property’s assessed value to its realized market value – within these tax jurisdictions. The average assessment ratio for a black resident in our sample is 12.7% higher than for a white resident. For black or Hispanic residents in aggregate, the average assessment gap is 9.8%. For the

same bundle of public services, minority residents are therefore paying a significantly larger effective property tax rate. For the median minority homeowner, the differential burden is an extra \$300–\$390 annually. This finding is strongly robust across most states in the U.S. We produce county-level estimates to characterize the distribution of this assessment gap. The average black homeowner in a county at the 90th percentile of the assessment gap distribution has a 27% higher assessment ratio, and would pay an extra \$790 annually in property tax.

We then explore two channels that drive these assessment gaps in the data. The first is spatial. We show that assessments are insufficiently sensitive to neighborhood-level attributes. Because of residential racial sorting, minority residents face, on average, different neighborhood characteristics than white residents ([Ananat 2011](#), [Cutler et al. 1999](#), [Massey and Denton 1993](#)). We show that assessed values and market prices align well on home-level characteristics, but diverge on tract-level attributes. In other words, market prices capitalize highly local factors, but assessments are much less responsive. This generates spatial variation in the assessment ratio within jurisdiction. Residential spatial sorting leads this variation to correlate with homeowner race and ethnicity, generating just over half of the average assessment gap.

The second channel is a racial differential that persists even after conditioning away spatial factors. Within U.S. Census block groups, which represent regions of approximately 1,200 people, an average minority homeowner has an assessment 5–6% higher relative to market price than her nonminority neighbor. This latter finding is particularly surprising given that most assessors likely neither know, nor observe, homeowner race. We document that a significant portion of this effect arises from racial differentials in assessment appeals. To do so, we first analyze appeals in Cook County, the second largest county in the U.S. Using administrative court records, we show that minority homeowners: (i) are less likely to appeal their assessment, (ii) conditional on appealing, are also less likely to win, and (iii) conditional on success, typically receive a smaller reduction than nonminority residents. Then we show that national assessment patterns around changes in racial ownership are consistent with this channel: within the same property, assessment growth is significantly higher during the tenure of a black or Hispanic homeowner.

Finally, we propose a solution to at least partially address racial disparities in assessment gaps. This solution is tractable, and only uses publicly available data. As explained above, our results suggest that at least half of the racial disparity in assessment ratios emanates from the failure of assessments to account for geographic variations in neighborhood attributes. We describe an algorithm for generating assessments that relies on small-geography home price indexes. We show that simply linking assessment growth to zip-code-level indexes will

reduce racial inequality by 55–70%. Racial inequality can be further reduced by using house price indexes that are more carefully calibrated to local geographies than zip code boundaries, which are well known to be drawn with little consideration for local characteristics.

We believe the results we uncover in this paper represent a large source of racial inequality in the United States. Property taxes are directly relevant to nearly everyone in America. Many of the most salient public goods including education, policing, transportation infrastructure, and utilities are provided chiefly by local governments: cities, counties, towns, school districts, and other special purpose entities. For most of these local budgets, property tax revenue is the central financial pillar. For the average local government, property tax receipts comprise 56% of general revenue; and for the 14,000 independent public school districts in America, the average proportion is even higher at 74%.¹ The jurisdictions we form represent regions where residents have at least tacitly agreed upon some level of intended taxation, and an associated level of public amenities provided with that revenue. Inequality within these jurisdictions, therefore, suggests that racial and ethnic minorities in the U.S. face different prices for the same set of public goods.

Much economic analysis of discrimination, including the canonical approaches of both Becker and Arrow/Phelps, focuses on how racial differences arise within a market environment. In contrast, our setting of local public finance allows us to study how racial differences in outcomes arise structurally, or institutionally, in a nonmarket environment. Since the Fair Housing Act of 1968, overt discrimination by race has been illegal. We show how inequalities arise nonetheless from institutional features of property tax administration. Our empirical analysis of assessment ratios is closely aligned to the legal concept of “disparate impact,” a term that denotes group-level differences in outcomes between protected classes, one of which is race.² While differences are permitted, average differences between groups delineated by any protected class constitutes a discriminatory outcome, regardless of the process which generates the disparity.³ Our results show that race-blind policies may still generate outcomes which are not race-neutral.

¹ General revenue excludes state and federal transfers, which can be large. General revenue is the funding stream which the local entity can directly affect.

² The others are: “religion, sex, handicap, familial status, or national origin” (42 USC 3604–3605). In 2015, the U.S. Supreme Court affirmed that disparate impact is the standard by which to legally evaluate discrimination claims in the housing market. *Texas Department of Housing and Community Affairs v. Inclusive Communities Project, Inc* (135 S. Ct. 2507). This remains an evolving area of jurisprudence. In August 2019 the Department of Housing and Urban Development issued a call for comments on a new proposed rule which would change the legal standards for establishing a disparate impact claim; this proposal is still outstanding as of this writing.

³ An exception is if the process is “necessary to achieve one or more substantial, legitimate, non-discriminatory interests” of the government ([Atuahene 2017](#)).

We make four main contributions to the literature. First, we contribute to the literature on racial disparities in the property tax. There is a long history of activism seeking to address racially motivated over-assessment of residential property. [Kahrl \(2016\)](#) describes property tax rates as central to African American political mobilization during the Reconstruction era, and also provides examples of homeowners in the 1920s and 1930s suing local governments for relief from discriminatory assessments. [Rothstein \(2017\)](#) details the same concerns arising in the 1960s and 1970s. [Atuahene and Berry \(2019\)](#) estimate a causal link between inflated assessments and tax foreclosures within one county in Michigan between 2009 and 2015.⁴ We build upon this research by: (i) documenting the widespread, contemporaneous presence of assessment gaps using comprehensive national data; (ii) providing a more refined notion of the proper taxing jurisdiction to precisely quantify the breadth and magnitude of disparate impact in property tax burden in the post-civil rights era; and (iii) evaluating the mechanism through which the assessment gap arises.

Second, we show a new mechanism which can help explain the large and persistent black-white wealth gap. One strand of the broad literature studying racial inequalities in wealth has focused on the role of geography and spatial sorting ([Cutler et al. 1999](#), [Gittleman and Wolff 2004](#), [Card and Rothstein 2007](#), [Charles and Guryan 2008](#), [Ananat 2011](#), [Chetty et al. 2014, 2018](#)). We show a new source of racial inequality in wealth, which operates through a public finance channel and is largely generated by spatial sorting within taxing jurisdiction. The effect is highly persistent (property taxes must be paid every year) and exists in most locations across the country. At the median, the assessment gap results in a black homeowner paying approximately \$390 dollars more each year. This is a very large number, given that median black household net worth is \$13,000, of which only \$4,000 is in liquid assets.⁵ For any discount rate below 3%, the stream of incremental tax payments suggested by our findings represents an excess tax burden which exceeds total household wealth for the median black family. Not only does this inhibit wealth building directly, it may well distort home ownership and financing choices for minority residents, further exacerbating the wealth gap. Much work in wealth inequality focuses on channels that affect education and wages; here we show a channel which operates on wealth directly.

Third, we contribute to a small but growing literature that explores the bias and distributional consequences of algorithms and statistical procedures. An active debate in this literature is whether using race or racially correlated variables will reduce or exacerbate bias in outcomes. [Bartlett et al. \(2018\)](#) show that FinTech algorithms in the mortgage market

⁴ In a related article [Atuahene \(2017\)](#) argues that present-day assessment practices in the city of Detroit should be considered federally illegal under the Fair Housing Act.

⁵ Survey of Income and Program Participation (2016).

generate higher interest rates for Hispanic and African American borrowers, although rejections are lower relative to face-to-face lending. Fuster et al. (2018) show that black and Hispanic borrowers are less likely to gain from increased precision in credit prediction generated by machine-learning models. Kleinberg et al. (2018) argue that allowing algorithms to use protected class variables directly will provide an effective mechanism for reducing bias in decision making. Our results in this setting support the latter notion. Automated valuation and mass appraisal is an algorithmic prediction problem. We show that assessments will more closely track market values if the demographic composition of local areas is considered, simply because this variable is a strong statistical proxy for many factors that influence market prices.

Finally, we add to an extensive literature examining racial differences across a diverse range of outcomes including health (Schulman et al. 1999, Williams 2012), employment (Donohue III and Heckman 1991, Card and Lemieux 1996, Bertrand and Mullainathan 2004), criminal justice (Knowles et al. 2001, Arnold et al. 2018), and residential housing markets (Charles and Hurst 2002, Bayer et al. 2007, Card et al. 2008, Bayer et al. 2017). That U.S. government policies in the early half of the twentieth century deliberately promoted racial segregation and discrimination in housing markets has been widely documented (Rothstein 2017); however overt discrimination by race has been illegal since the 1960s. We use modern data to show that minority homeowners still face financial discrimination generated at the intersection of housing markets and local public institutions.

The paper proceeds as follows. Section 2 describes the typical structure of local property taxation and highlights important institutional details that pose econometric challenges. Section 3 outlines our empirical strategy. Section 4 details the data sets we use. Section 5 presents the results. Section 6 outlines potential policy approaches for achieving a equitable tax burden. Section 7 concludes.

2 Setting and Institutional Detail

2.1 Local Governments

Property taxation in the United States is chiefly a feature of local governments. Government authority in the United States is organized at three levels: federal, state and local.⁶ These levels are roughly hierarchical. State constitutions and laws empower local units of govern-

⁶ The intent of this overview is to help orient any reader unfamiliar with the general structure of American government. It is very much not a careful description of American federalism or the ways in which authority is mediated between levels of government.

ments, while retaining preeminence in the case of any regulatory conflict. Local units are empowered either by an explicit enumeration of powers,⁷ or through “home rule” provisions which grant local units all authority not explicitly reserved for the state.⁸ Counties and cities are the most prominent example of local governments, though as discussed below there are many other relevant types of local entities.

Although lowest in the hierarchy, state and local governments tend to have the most salient day-to-day impact on the lives of residents. Schooling, public safety (police and fire), infrastructure, and transportation are all amenities that are chiefly provided by local governments with varying degrees of state and federal support. The vast majority of local government units impose a property tax, and these revenues are the central fiscal pillar of local government budgets.⁹ They are the largest source other than intergovernmental transfers, and are very stable year-over-year (see Table 1). For some important local amenities, property taxes are even more crucial: at 74%, independent school districts are almost entirely dependent upon property tax revenue. In 2012, local units of governments collected an aggregate of \$433B in property taxes.¹⁰

Local units have broad discretion to set the level of intended tax burden. There is wide regional heterogeneity in the mechanism used to change the tax rate. Two approaches are most common: either voters have direct input into property tax levels at the ballot box, or they delegate this authority to locally elected officials (who may or may not redelegate this authority to appointed individuals). Often, the intended tax burden is implicitly defined: a certain level of spending will be approved (through either of the previously mentioned mechanisms) and then this amount will be divided by the total value of the local property, yielding an implicit tax rate. For this analysis, what is important is that local units set their own intended level of taxation each year.¹¹

⁷ Oklahoma is one example. See Article XVIII–1 of the State Constitution and the extensive codification of authority in Title 11 of the Oklahoma statutes.

⁸ Montana is one example. Typical language appears in Montana Code Annotated 2019, Title 7, Ch 1, Part 1.101, and Article XI, Part XI, Section 6 of the state Constitution.

⁹ Some states impose a state-wide property tax levy, but the major source of state revenues are sales and income taxes. In 2012, state governments accounted for only 3% of all property taxes raised; local units comprised the remainder.

¹⁰ Authors’ calculations using Census of Governments data; figure given is nominal dollars.

¹¹ There are often legislative constraints that limit the rate of annual change.

2.2 Effective Rates Depend on Assessments

The residential property tax is implemented as an ad valorem tax: residents are taxed some proportion of the value of their property.¹² This concept is regularly explicitly delineated in American law. Virtually every state has language in its constitution or legislative code carefully specifying that property taxation is intended to represent a proportional burden on the fair market value of the real property.¹³

As described in Section 2.1, for every taxing government, a tax rate exists (either explicitly or implicitly defined). We refer to this as a "policy rate" to highlight that this rate is a political or legislative lever which can be adjusted to change the desired tax burden. However, it is crucial to realize that this policy rate is, alone, not sufficient to characterize the effective tax rate. This is because tax bills are calculated by applying the policy rate to an "assessment," which is a local official's projection of a home's value. For every home in America, there is some bureaucratic entity charged with producing an assessed value for that property. Very often – but not always – this responsibility lies with county governments, and is executed through a county assessor's office. These property assessments are a legal determination of value for purposes of the taxing entity, and will be a central object of our analysis. Assessments are produced for every home in a jurisdiction. Assessments are revised annually, sometimes biannually, or in some regions even less often.¹⁴ We observe realized assessments for all homes in our dataset annually.

If the policy rate is 5% and the home's assessed value is \$100,000, then the homeowner will receive a tax bill of \$5,000: a 5% tax applied to the \$100,000 assessment. However, and perhaps surprisingly, nothing in the previous sentence necessarily implies that the market value of the home is \$100,000. While the natural intuition might be to assume that assessed values track market values one-to-one, this is not the case for most of the country. Local units have a free scaling parameter in choosing how to produce assessments. States may mandate a particular level: Alabama specifies that residential assessments should be 10% of market value.¹⁵ Thus, if the home described in the beginning of this paragraph were in

¹² While there are examples of localities imposing fixed, or unit property taxes, these tend to be specific levies approved to fund a particular project (or to cover debt service for a given bond issuance). We do not have any way of providing an aggregate breakdown of tax dollars raised by ad valorem taxation versus unit taxes, in every region we have looked at specifically, unit taxation is a very small portion of overall proceeds.

¹³ One example from Georgia: "'Fair market value of property' means the amount a knowledgeable buyer would pay for the property and a willing seller would accept for the property at an arm's length, bona fide sale." 2018 Georgia Code, Title 48, Chapter 5, Article 1(3).

¹⁴ Cook County, IL for example, conducts assessments on a triannual scale; each property is assessed every third year.

¹⁵ Code of Alabama, Section 40, Chapter 8, Section 1.

Alabama, the \$100,000 assessment would, in fact, imply a market value of \$1,000,000.¹⁶ The effective rate of taxation for this home would be .5%: the homeowner pays \$5,000 in tax on a million-dollar asset.

In absence of state regulation, local units choose their own scaling factor. Sometimes local practices conflict with state targets, adding another layer of administrative complexity: Illinois state dictates that assessments should be 33%, but Cook County within Illinois uses 10% as an assessment target. To reconcile these, Illinois state law mandates an adjustment be applied to local assessments in order to achieve the state-level target. Figure 2 shows the raw distribution of realized scaling factors in our data. As this figure shows, there are jurisdictions which appear to be targeting 100%, but there are also many jurisdictions which are clearly targeting another number.

Economically speaking, assessment scaling factors are meaningless. It is simply a free choice parameter for the local government. Consider tax revenue in jurisdiction j and year t as a function of two variables: $revenue_{j,t} = f(rate_{j,t}, scaling_{j,t})$. Revenue is stable if the rate is doubled when the scaling factor is halved. However, as a consequence of the regional heterogeneity in scaling, the econometrician observing only two homes, each assessed at \$50,000, can make no inference about relative market value of these properties and hence also cannot discern whether the tax burden for these properties should be the same.

We do not observe scaling factors. In order to draw meaningful inference from variation in assessment ratios, we must be able to discern how all properties map geographically to governments. The following section outlines the challenges this poses, and describes our solution, which is to form “taxing jurisdictions” that hold fixed the choice of scaling factor, as well as the level of intended taxation.

2.3 Forming Taxing Jurisdictions

Local governments are highly spatially complex. Across the U.S. more than 75,000 entities potentially impose a property tax. Homeowners typically face taxation from multiple local units simultaneously. As mentioned, cities and counties are key examples of local government units. However, it is very common for regions to have a range of separate autonomous taxing entities. Chief examples here are: school districts, park districts, and municipal utility districts. Taxing authority may also be embedded in a special purpose district like an airport authority or regional economic development initiative. As a rule, the boundaries of these units are not naturally coincident. Counties are a complete partition of space in the

¹⁶ Errors in assessed values are central to this paper. This calculation assumes an accurate assessment for purposes of the example.

US: every point in a given state lies in exactly one county. However, no such logical precision applies to other local entities. Cities often lie across county boundaries. In low-population-density areas, school districts often cover multiple towns (and potentially portions of different counties); in urban areas, there may be multiple school districts within a given metropolitan region. Units like park districts or utility districts typically have a delineation governed by a service area that reflects physical geography and may have little to do with nearby civic boundaries. Excluding state governments, the average home in the United States is touched by 4.5 local entities, all of which potentially levy a property tax.¹⁷

Our empirical goal is to explore whether minority residents face a higher tax burden than their white neighbors, *conditional on holding intended taxation fixed*. We are not asking whether minority residents tend to live in regions with more (or less) public services, which would naturally suggest higher (or lower) taxes. Rather, we wish to compare two residents for whom the tax burden should be identical: served by the same set of governments, receiving the same bundle of public goods, and facing the same policy tax rate. This analysis is only possible because we find a method for discerning the networks of overlapping governments that touch properties in our sample. We accomplish this by mapping the geolocation of property parcels onto geographic shapefiles for the universe of local governments in the United States. Every piece of empirical evidence in this paper rests on holding these unique government networks fixed. Colloquially, we will use the term “jurisdiction” to denote the set of unique taxing entities that touch a given property.

Panel A of Figure 3 illustrates this approach in a stylized example. There are three governments in this example: the county, which contains a city and an independent school district. The city and the school district have partial overlap. This spatial overlay of governments generates 4 taxing jurisdictions. Jurisdiction one contains those homes which receive services from, and are taxed by, the county alone. Homes in jurisdiction two are served and taxed by both the county and the city. Homes in jurisdiction three are served and taxed by all governments, and homes in jurisdiction four are served and taxed by the school district and the county. Panel B of Figure 3 highlights our focus on within-jurisdiction inequality. For any cross-jurisdiction comparisons, we cannot rule out Tiebout sorting along preferences for public goods or intended levels of property tax. Our focus is solely on inequality between residents who are subject to the same set of taxes and who have access to the same bundle of public goods.

The example in Panel A of Figure 3 is, in fact, quite common across the country. However, jurisdictions can be complex, especially in more urban regions. Figure 4 shows the example

¹⁷ Author's calculations using Atlas Muni Data shapefiles.

of Harris County, Texas. Including the county, there are 12 local units of government which overlap in varying combinations. Each combination forms a distinct jurisdiction. One such jurisdiction is the region defined by the nexus of all 12 governments (this region is not visually identifiable in Figure 4). In our full sample, we observe a market transaction (paired with an assessment) for approximately 100 homes within this particular jurisdiction. This is a relatively small jurisdiction. Others are the size of cities and encompass tens of thousands of home transactions.¹⁸

The within-jurisdiction analysis is crucial in two distinct ways: (i) it holds fixed the level of intended tax burden, and (ii) it holds fixed the regionally chosen scaling factor. As we will describe and support, equity in the property tax demands that assessment ratios be constant within jurisdiction. This is not a controversial notion: it is often mandated in state legal codes,¹⁹ and is also a primary tenet of best-practice standards for professional assessors.²⁰

Assessment ratios, therefore, are only relevant because they are a sufficient statistic for inequality in effective tax rates. This logical relationship only holds, however, within a region where everyone should be facing the same tax burden. No meaningful comparison in tax burden is possible if we compare residents who pay taxes to (and thereby receive services from) a different set of governments. The way we form jurisdictions ensures that we only compare tax burden between residents paying tax to the same set of governments.

Similarly, we cannot meaningfully compare assessment ratios between two homeowners who live in regions which are simply targeting a different assessment ratio. Our data does not reflect which entity produces assessed values, nor the target assessment ratio. Conducting our analysis within jurisdiction—according to our precise notion thereof—ensures that no error is introduced by an inability to observe local unit heterogeneity in assessment practices.

While our jurisdictions have both a natural economic and political interpretation, it is certainly reasonable to wonder whether our results are driven in any way by the partitioning of geography. We can test this fairly directly. Practically speaking, assessments are most commonly done at the county level. Often this is a provision of state law, but even when not required, it seems that either custom or natural considerations of efficiency and resource management often result in counties “owning” assessments. While it does not make any sense to compare effective tax rate within county (because so many sub-county units impose other property taxes and provide services), if target assessment ratios are unlikely to vary within

¹⁸ In some regions, all substate units of government are spatially aligned; Philadelphia is one such example: the county and city of Philadelphia, along with the school system, are all entirely coincident. This is relatively rare.

¹⁹ See for example, Michigan Compiled Laws, Section 211.34(2).

²⁰ For example: *Guidance on International Mass Appraisal and Related Tax Policy*, International Association of Assessing Officers, 2014.

county, we can meaningfully compare assessment ratios within county instead of within jurisdiction. In Section 5, we show that our baseline results establishing racial differences in assessment ratio are robust to conducting our analysis within county. Within county estimations, in fact, generate slightly higher estimates. Our preferred specifications all use the more rigorous partitioning into jurisdictions of unique overlapping governments.

3 Empirical Strategy

The outline of our approach is as follows. We first define a notion of equitable tax administration within a jurisdiction. Then we show that within-jurisdiction variation in assessment ratios is an empirical sufficient statistic for rejecting the equitable tax null.

Our notion of equity relies on the ad valorem nature of the property tax, and the fact that taxes are levied on assessed values. Let i denote property, j taxing jurisdiction, and t year. Further, let V^* be the true value of the property being taxed. Given a tax rate of $r_{j,t}^{eff}$, by definition an equitable ad valorem tax must satisfy:

$$\text{equitable tax}_{i,j,t} = r_{j,t}^{eff} V_{i,j,t}^* \quad (1)$$

Note that r^{eff} is an effective tax rate. Let c be the local scaling factor for assessments, and let r be the policy rate that rationalizes equation 1: $r_{j,t}^{eff} = r_{j,t} c_{j,t}$. This last equation simply reflects that if assessments are scaled to be half of market value, the policy rate must double in order to achieve the level of tax burden implied by r^{eff} .

Property tax bills are actually generated by applying this policy rate to an assessed valuation, $A_{i,j,t}$:

$$\text{actual tax}_{i,j,t} = r_{j,t} A_{i,j,t} \quad (2)$$

Our equitable tax null is simply that $\text{actual tax}_{i,j,t} = \text{equitable tax}_{i,j,t}$. We observe $A_{i,j,t}$, the realized assessed valuation assigned to the house. We observe market prices for homes, $M_{i,j,t}$, and accordingly will let $M_{i,j,t} = V_{i,j,t}^*$.²¹ Equating 1 and 2, and taking logs yields a necessary condition for equitable administration of an ad valorem tax:

$$\ln(A_{i,j,t}) - \ln(M_{i,j,t}) = \ln(c_{j,t}) := \gamma_{j,t} \quad \forall i \quad (3)$$

²¹ It is worth reiterating that state laws explicitly state that property taxation should be levied upon the “fair cash value” that would be received in an arm’s-length transaction. Therefore, our reliance on market prices is not a strong statement about perfect markets or market efficiency, but rather a reflection of the legal intent underlying the taxation.

Equation 3 is a theoretical statement that does not allow any errors at all in assessments.

Empirically, we define a deviation from our fair tax benchmark in context of arbitrary demographic delineations. Partition the homes of any jurisdiction into M subsets based on any demographic characteristic, and denote by $m \in \{1, 2, \dots, M\}$. Let $\bar{c}_{m,j,t} := \frac{1}{N} \sum_{i \in m} c_{i,j,t}$. Our fair taxation null is:

$$\bar{c}_{m,j,t} = \bar{c}_{m',j,t} \quad \forall m, m' \quad (4)$$

Equation 3 says that assessment ratios should not vary at all within jurisdiction. While strictly true, this represents unattainable precision. Equation 4 says that average assessment ratios should not vary within jurisdiction by demographic group. For an ad valorem tax burden within a jurisdiction, taxes levied should be a constant proportion of market value. However, property taxes are calculated as a proportion of assessed value. Thus, an increase of assessments relative to market value represents an increase in the overall tax burden. Correspondingly, if group m has higher assessed valuations relative to market than group m' , then group m faces a higher tax burden.

We test inequality by racial and ethnic group with estimating equation:

$$\ln(A_{i,j,t}) - \ln(M_{i,j,t}) = \gamma_{j,t} + \beta^r race_{i,j,t} + \epsilon_{i,j,t} \quad (5)$$

Here $race$ is a vector of indicator variables for racial and ethnic groups. The formulation of equation 5 is motivated by the legal notion of disparate impact. The Department of Housing and Urban Development states: “[a] practice has a discriminatory effect where it actually or predictably results in a disparate impact on a group of persons [...] because of race, color, religion, sex, handicap, familial status, or national origin.”²² As the left-hand side of equation 5 is a sufficient statistic for within-jurisdiction tax burden, this formulation is of primary interest for establishing inequality by race. The fixed effect $\gamma_{j,t}$ absorbs the realized average assessment ratio within jurisdiction. Then, since $race$ is a categorical variable, β^r is a vector of estimated group-level deviations from average realized assessment ratio. If β^W , the average assessment ratio for white residents, is statistically different from β^M , the average assessment ratio for any grouping of minority residents, this would be evidence of inequality in tax burden.

Finally, before proceeding, we highlight an econometric point. Given the way tax bills are generated, our equitable benchmark implies a constant simple assessment ratio. We use logged values in our specifications because of heterogeneity in scaling factors. As we show in Figure 2, some regions target an assessment ratio below 10%; others 100%, with a

²² 24 CFR 100.500(a).

wide range in between. Therefore an aggregate regression using simple ratios, which implies additive rather than proportional deviations around the mean ratio, does not have a natural interpretation. Scale invariance makes log differences a natural solution, however, this does generate a Jensen's inequality term. If minority residents select into homes with higher within-jurisdiction variation than white residents, this would bias our estimates upwards. In Section 5 we show that our results are robust to running regressions on simple ratios rather than log differences.

4 Data

The core research design of this paper rests on combining data from three sources: 1) property-level panel records of assessments, transactions, home characteristics and geolocation from ATTOM, 2) Geographic Information System (GIS) detail on local government boundaries from Atlas Muni Data, and 3) mortgage-holder race from Home Mortgage Disclosure Act records. These three sources are merged to create a panel of observations at the property-year level. For each home, four pieces of information are observed: (i) the network of taxing entities touching that property, (ii) the annual assessment, (iii) whether any transaction occurs, along with the transacted price if so, and (iv) the race of the homeowner (both buyer and seller in the case of a transaction). For any analysis of assessment ratio, we restrict attention to homes which transact in an arms-length sale with an observed market price, and we focus on the race and ethnicity of the home seller (the homeowner at the time when the assessment was done). We merge this assembled dataset with standard data from the U.S. census and the American Community Survey.

One salient choice we make is to remove all California properties from the final dataset. While taxation in California is legally characterized as an ad valorem tax, the state passed Proposition 13 in 1978, amending the state constitution to place extremely stringent limitations on assessment practices. Assessment growth within each homeowner's tenure is capped at 2% annually, which is far below the growth in market prices in almost every region. In addition, although assessments are supposed to revert to market value upon sale, Proposition 13 also provides a range of mechanisms by which groups of homeowners can transfer the artificially low basis from one property to another. The statewide result has been a decades-long divergence between assessments and market values. As a result, we consider the property tax in California to be ad valorem in name only.²³ Our analysis, applied to the

²³ Nonetheless, the constitutional amendment authorized by Proposition 13 continues to describe ad valorem property taxation, as in e.g. California Constitution, Article XIII.A Section 1(a).

California data, does show similar evidence of racial and ethnic inequality. For completeness, we show this in Section 5.1. However, our subsequent analysis of mechanisms in this paper is less relevant for California, simply because assessments there are so mechanically driven by the restrictions of Proposition 13.

4.1 Property Records

We obtain property-level records of assessments and transactions from ATTOM. This is a comprehensive dataset with annual observations on 118 million properties in the U.S. from 2003–2016. Assessment and transaction records are sourced from county assessor and recorder offices respectively. Each property is characterized by a unique identifying ID, which allows us to match assessments with transactions. In addition, each property has a use code which ATTOM harmonizes across local definitions. We restrict our attention to residential properties of up to four units. Commercial property is generally assessed differently from residential properties, so we cannot draw inference from jurisdiction average assessment ratios without restricting to residential properties only. Further, multi-family homes (e.g. large apartment buildings) are sometimes subject to different assessment rules. The restriction to residential properties of one to four units gives us a set of properties which should always be assessed in the same way within jurisdiction. In order to avoid having to impute any market values, our baseline dataset includes only homes for which we observe the sale price in an arm’s-length, full consideration transaction. The recorder portion of the ATTOM dataset has several indicator flags for arm’s-length transactions and partial interest sales, which collectively can be used to isolate transactions that reflect an accurate signal of market value. The ATTOM data also provides a record of tax dollars paid by the homeowner, along with any exemptions. Importantly, each home is identified with a latitude and longitude for the parcel. These are used to geolocate the home within government borders.

4.2 Government Boundaries

We obtain shapefiles for government boundaries from Atlas Investment Research’s Atlas Muni Data. These 75,000 shapefiles are intended to span the universe of local governments in the U.S. The core set of shapefiles covers counties, cities, towns, schools and special districts as defined by the U.S. Census. In addition, Atlas Muni Data developed proprietary shapefiles for any entity which has ever accessed public markets, as compiled from Municipal Securities Rulemaking Board filings. Thus, a shapefile is developed for any entity which has ever issued either general obligation or revenue bonds. As debt issuance is very often paired with either broad authority to tax (in the case of general obligation bonds) or a voter-

approved one-off tax levy (more common for revenue bonds), we consider each of these entities as a potential taxing entity. Collectively, in addition to the 50 states and D.C., the Atlas data covers 3,142 counties, 46,660 cities or towns, 13,709 independent school districts, and 11,924 special purpose districts.

We use all of these shapefiles to form our taxing jurisdictions. This is a very robust and flexible empirical strategy: if any given entity does not tax, we do not introduce any bias by considering it in forming our unique government networks. If anything, we create another barrier against observing any distortion by restricting our analysis to a (potentially) smaller geographic region. And, of course, if the entity does levy a tax or do its own assessments then failing to take it into consideration would certainly introduce bias.²⁴ Each shapefile delineates a region in space by connecting a large number of latitude and longitude segments. We use standard GIS techniques to associate each home's longitude and latitude with any shapefile that contains that point. It is an embedded assumption that the latitude and longitude of the property correctly characterizes government association.

We place emphasis on the comprehensive nature of these shapefiles. One significant threat to our research design would be an inability to observe any assessing entity. Practically speaking, because this function so often is assigned to counties or large cities, and because we have shapefiles for every county, and essentially every city and town, we feel that it is very unlikely we have missed such an entity. This is a strength of demonstrating distortions by using assessment ratios. Any statement that we make about tax dollars also has a source of error if we have missed any taxing entity. The breadth of the government shapefiles suggest that any taxing entity not captured in the data is likely to be small.

4.3 Home Mortgage Disclosure Act Records

The Home Mortgage Disclosure Act (HMDA) mandates that financial institutions disclose certain information about mortgage applications and mortgage origination at an individual loan level. This law was enacted to provide transparency about credit access for minority residents and within historically redlined neighborhoods. One requirement of the law, therefore, is for financial institutions to solicit and report the racial and ethnic identity of loan applicants. Clients are asked their race and ethnicity directly; the designations are the same as the U.S. Census. A customer can decline to provide this information, and a missing flag is reported as well.²⁵ HMDA applies to financial institutions meeting certain criteria—the major

²⁴ We suspect it is rare for entities other than counties, cities, or towns to produce assessments; but our strategy is robust to such an instance.

²⁵ Regulation C of HMDA also requires loan officers to note race and ethnicity race based on visual observation if the application is made in person, and the applicant does not provide the information. During

one being an asset threshold which is currently \$46M for depository institutions and \$10M for for-profit mortgage lenders. During the 2005–2016 period we consider, between 6,900 and 8,900 institutions reported loans ranging in number from 14.3M to 33.6M annually.²⁶

We merge the HMDA records to the ATTOM dataset. This is a fairly standard merge in the literature (see, e.g. [Bayer et al. 2017](#) or [Bartlett et al. 2018](#)). HMDA loan records are uniquely identified by: year, census tract, lender name, and dollar amount (rounded to thousands). The ATTOM data contains: transaction date, latitude and longitude of the property, lender name, and dollar amount. We restrict our sample to the highest quality matches, requiring an exact match on year (permitting a one-month overlap between December and January), an exact match on tract, an exact match on (rounded) transaction amount, and a fuzzy string match on lender name.²⁷

The initial merge establishes race/ethnicity of the mortgage-holder at the transaction date.²⁸ Our end goal is to relate assessment ratios to homeowner race and ethnicity at the time when the assessment was generated. Assessments are produced in advance of the tax year in which they will apply.²⁹ Therefore, we exploit the dynamic structure of the transactions dataset to build a panel of homes for which we know the declared race and ethnicity of the homeowner at each year. There are two relevant cases: (i) sales and (ii) refinance transactions. For sales, the transaction pins down the race/ethnicity of the buyer, which is then associated with that property in each subsequent year until the next observed transaction. For refinance transactions, we carry race and ethnicity not only forward in time, but also backward, as the home does not change ownership. For a large number of transactions, race/ethnicity is not observed.³⁰ In these cases, we mark race/ethnicity as unknown, and carry that categorization forward and backward in time as appropriate. We fill our panel in this manner, with racial and ethnic indicators updating each time we observe a transaction. As a last step, we remove the observations for which mortgage-holder race and

the period covered by this paper, financial institutions were not required to distinguish between application-disclosed information and visually-observed information in reported data.

²⁶ Summary statistics from www.ffiec.gov.

²⁷ The diversity of retail-outlet names within a single financial institution can make exact string-matching a challenge in some regions. We rely on a natural language algorithm developed by the Real Estate and Financial Markets Laboratory at the Fisher Center for Real Estate and Urban Economics to match names. The algorithm trains itself within region on perfect singleton matches across all variables other than name, and then uses that mapping to assign a confidence index to each HMDA-ATTOM string-pairing.

²⁸ HMDA records also include information on co-applicants. We use race and ethnicity of the primary applicant only.

²⁹ In Philadelphia, for instance, the tax year runs from April 1st to March 31st, with payments due by March 31st. For the 2020 tax year, the Office of Property Assessment sent residents assessments at the beginning of April 2019.

³⁰ This occurs if we cannot match the transaction to a record in HMDA – in the case of a cash transaction, for instance; or if the race/ethnicity is recorded as “not provided” in HMDA.

ethnicity is unknown or not declared. We also remove any home which sells in consecutive years. This is because we do not perceive the exact timing of assessment generation. The approach described associates the assessment ratio from a transaction at time t with the race and ethnicity of the homeowner at time $t - 1$. Therefore if there are multiple homeowners during year $t - 1$, we cannot be sure how to assign race and ethnicity.

Our final baseline dataset is a panel of 6.9M homes. The data are anonymized: each home is characterized by a unique ID variable. For each observation, we have an assessment ratio based on a market transaction, know the associated taxing jurisdiction, and have the reported race and ethnicity of the homeowner. Each home is associated with a census tract and a census Block Group, permitting us to merge in a range of tract-level variables from the American Community Survey five-year estimates.

5 Results

The main results of our analysis are organized into five parts.

We first establish the existence and magnitude of the assessment gap in section 5.1. These results document the additional property tax burden faced by an average minority citizen. The analysis is within taxing jurisdiction, which ensures that we are comparing residents who: (i) face the same intended level of taxation, and (ii) are served by the same set of public institutions and government entities. To characterize the distribution of the assessment gap, we present state-level and county-level estimates. We also show that the average assessment gap is increasing in county-level minority population share.

Our second set of results, described in section 5.2, decomposes the assessment gap into two channels. The decomposition is along spatial lines. One channel, which we term “neighborhood composition,” relates to spatial variation in the assessment ratio, and operates through characteristics of a home’s geographic surroundings. Even within jurisdiction, people of different races live—on average—in different types of areas. This residential spatial sorting by race is very well known (Bayer and McMillan 2005, Logan and Parman 2017, Licher et al. 2007, among many others). We use hedonic regressions to show that market prices and assessed values are well aligned on the valuation of property-level attributes. However, there is large misalignment on pricing of tract-level variables. The magnitude of market-hedonic prices tends to be substantially larger than assessment-hedonic prices. This suggests that assessors are insufficiently taking neighborhood factors into account when constructing assessments. Given the pattern of racial sorting that exists, the result is assessments that are on average too high relative to market prices in the regions that have a higher minority population share.

The other component of the assessment gap is a racial differential that persists even after conditioning away spatial factors. We establish this finding by conducting our analysis within small geographic units to control for neighborhood-level variation. We are implicitly comparing two homeowners of differing race within the same census tract (approximately 4,000 residents) or census block group (approximately 1,200 residents). We refer to this as a “homeowner effect,” and posit that one mechanism relates to assessment appeals.

Our third set of results presents evidence on the role of assessment appeals in generating inequality in the property tax. To the best of our knowledge, there is no national compiled dataset on property assessment appeals. In section 5.3, we test the role of assessment appeals in generating inequality, using administrative microdata spanning 12 years of assessment appeals from the second largest county in the U.S. We show racial and ethnic differentials in appeals behavior and outcomes, even within tract and block group. Then we show that assessment patterns nationally are consistent with the appeals channel that we document in a single county.

In section 5.4, we analyze heterogeneity in the assessment gap by racial attitudes and regional minority population share. We first use the measure of racial animus described in [Stephens-Davidowitz \(2014\)](#) to see whether the assessment gap varies with regional racial prejudice. The assessment gap is larger in areas with above-median animus, but is also large and statistically significant in below-median areas as well. This holds both for the overall estimates, and the homeowner effect estimates. We also split our sample into quintiles by average county-level minority population share, and show that the assessment gap is increasing in minority share.

In section 5.5, our fifth set of results shows that assessment gaps do, in fact, lead to higher tax burdens upon minority residents. While this is the natural implication of assessment ratio distortions – indeed the tight link between assessment ratios and effective tax rates is precisely what motivates our focus on the ratio – we close the loop empirically by demonstrating that this link does hold in the data.

We conclude the Results section with additional discussion of two points: (i) the central role of market prices in our empirical strategy, and (ii) the interplay between income or wealth and the assessment gap.

5.1 Assessment Gap Baseline

As outlined in Section 2, assessment ratios should be constant within jurisdiction for all residents. As a theoretical statement, this is a necessary condition for an equitable tax benchmark. In practice, the average assessment for any arbitrary grouping of residents

should be statistically equal. If the groups in question are distinguished by race (or any other protected class), different group averages represents a discriminatory outcome.

We establish our benchmark finding of an assessment gap by showing that assessment ratios within jurisdiction are, in fact, higher for minority residents. Following equation 5, we regress assessment ratio directly on a categorical variable for racial and ethnic groups, along with a jurisdiction-year fixed effect to absorb variation arising from regional scaling choice. Our equitable tax null implies a statistical zero for any race or ethnicity covariate.

Across all our results, we consider three groupings of minority residents. One is mortgage holders whose racial identification in HMDA is “black or African American.” The second combines the two largest racial and ethnic minorities in the county: anyone whose racial identification in HMDA is “black or African American” or whose ethnic identification is “Hispanic or Latino.”³¹ The third is mortgage holders identified in HMDA as having any race other than white or black, and not of Hispanic or Latino ethnicity. This last grouping is not a natural division and masks a large amount of underlying racial heterogeneity. The data is not sufficient to conduct a more precise racial breakdown, or a county-of-origin breakdown. We include these results for the sake of completeness. In all cases, the comparison group is non-Hispanic white residents.

Table 1 presents our baseline results. Within jurisdiction, assessment ratios are 12.7% higher for black homeowners, 9.8% higher for black or Hispanic homeowners, and just under 3% higher for other nonwhite homeowners. Given a national median effective property tax rate of approximately 1.4%, and a median home value of approximately \$207,000, this translates to an additional tax burden of \$300–\$390 per year.³²

We show two results that characterize the distribution of the assessment gap. First, Figure 5 shows the assessment gap by state for black residents and for black and Hispanic residents. We present results only from states with at least 500 observations, which excludes seven states. In the remaining set, the assessment gap is positive and strongly statistically significant in most states. For black homeowners, the state level estimates range from 33% to -3%. Estimates are positive and significant in 34 states, positive and insignificant in 5, and negative and insignificant in 3. For black or Hispanic homeowners, the pattern is very

³¹ HMDA regulations do permit loan officers to record race and ethnicity on the basis of visual observation, rather than soliciting the information from the client. Therefore, while this information is often self-reported, that is not always the case.

³² Averaging over white, non-Hispanic residents, the median jurisdiction in our data realizes an effective tax rate of 1.4%. Other methods of computing a national median property tax rate return similar figures. We obtain a median home value of \$207,000 for minority homeowners by taking Zillow’s national 2019 estimate of \$231,000, and reducing it by 10%, which reflects the ratio of black or Hispanic-owned home value to median home value in our baseline dataset for the latest available year (2016).

similar.

Second, we estimate the assessment gap at a county level. The results for black residents are shown in Figure 6. The distribution for black and Hispanic residents grouped together has a very similar shape. We again restrict attention to counties which have at least 500 observed assessment ratios. This reduces our sample to 671 counties. Our estimates range from 54% to -49%. The interquartile range is 14.8% to 4.7%. Point estimates are positive and significant at the 5% level in 391 counties, positive and insignificant in 219 counties, negative and insignificant in 53 counties, and negative and significant at the 5% level in 8 counties. For a black homeowner at the 90th percentile of this distribution, the assessment gap would be 27%. Again considering a \$207,000 home subject to a 1.4% tax rate, this would translate into an additional tax burden of \$790 annually.

As discussed in Section 4, we exclude California from our main analysis. It is widely known that California's property tax has been significantly distorted by Proposition 13, which caps assessment growth at 2% annually during any homeowner's tenure. Under particular circumstances, homeowners can also carry these artificially low assessments from property to property, and can bequeath them to their immediate heirs. In California, for most locations during our sample period, the growth of market prices was significantly greater than 2% annually. For this reason, the primary driver of inequality in the California property tax is more likely related to differentials in homeowner tenure and regional price appreciation, rather than the mechanisms we explore below. For completeness, in Appendix Table A1 we show the results of our baseline analysis applied to California. We do not present other results related to California in this paper.

In Appendix Table A2, we re-estimate the assessment gap using county-year fixed effects rather than jurisdiction-year. The point of this exercise is to show that our careful partitioning of space into taxing jurisdictions is not somehow mechanically driving our results. Differing levels of intended taxation by cities, towns, schools and others makes a within-county analysis of effective tax rate meaningless. However, counties are most often the entity which produces assessments. We can therefore reasonably consider assessment ratio variation within county-year. The results are very consistent with our baseline finding. Inequality in assessment ratios is approximately 4% higher within-county than it is within-jurisdiction. Our preferred specifications all employ the more rigorous within-jurisdiction analysis, not only because it is more likely to hold local assessment practices fixed, but more importantly because jurisdictions are able to hold fixed intended level of taxation and the set of entities providing public services.

Appendix Table A3 shows that our results are unaffected by using jurisdiction-month-year fixed effects instead of jurisdiction-year fixed effects. Municipal entities often employ a

fiscal year that begins partway through the calendar year (July and October are particularly common starting months). Ideally, our jurisdiction-year fixed effects would align exactly with the fiscal cycle selected by local taxing units, in order to absorb the effect of any deliberate change in assessment practices between fiscal years. We do not observe the local choice of starting month for the fiscal year. Appendix Table A3 shows that our estimates are robust to looking within jurisdiction-month-year. This suggests that any error introduced by forming fixed effects using calendar years rather than the (infeasible) fiscal years does not meaningfully change our estimates.

Finally, we return to the econometric point discussed at the end of section 3. Given large regional heterogeneity in the scaling factor, we use log assessment ratios in all regressions. This does potentially introduce bias in the form of a Jensen’s inequality term. Appendix table A4 we show that our baseline findings are robust to using simple ratios as the dependent variable. The estimates are naturally lower, as the coefficients are weighted averages of variation around target ratios ranging from 7% to at least 100%. We provide this table only to show that our use of log differences on the left hand side in equation 5 is not mechanically driving our results; the estimates themselves have no natural interpretation.

5.2 Neighborhood Composition and Homeowner Effect

We show that the assessment gap arises in two distinct ways. In the preceding section, we show variation in assessment ratio within taxing jurisdiction that correlates with homeowner race. Because these regions are carefully constructed to hold taxing units and policy rates fixed, variation in the assessment ratio represents a deviation from an equitable tax benchmark.

Within jurisdiction, there is also a large amount of spatial variation in assessment ratio, and this covaries in striking ways with race. Figures 1–4 show this spatial component in two large counties: Philadelphia County in Pennsylvania, which is coextensive with the city of Philadelphia; and Cook County in Illinois, which contains most of Chicago and several surrounding suburbs. Figure 7 is a demographic heatmap of Philadelphia at the census tract level. Using the American Community Survey five-year estimates, we plot the share of black or Hispanic residents in each tract. Figure 8 shows within-jurisdiction variation in realized assessment ratios. If the property tax were equitable in Philadelphia, the map in Figure 8 would be all the same color. Clearly this is not the case. In addition, there is very high spatial correlation between assessment ratio and minority population share. Figures 9 and 10 provide a parallel view of Cook County.³³

³³ Cook County touches numerous towns and other taxing units, and therefore contains multiple juris-

These spatial patterns are present to varying extent in many counties and cities. Our decomposition of the assessment gap will disentangle this spatial and geographic (between-neighborhood) variation from the non-spatial (within-neighborhood) drivers. We proceed by establishing the magnitude of the assessment gap while holding neighborhood attributes fixed. We show this is 45–50% of the total assessment gap. The remaining 50–55% then is the between-neighborhood variation. We explore this further in section 5.2.2.

5.2.1 Homeowner Effect

We proceed by showing that estimates of racial differentials stabilize as we condition on smaller and smaller geography. This exercise approximates the ideal experiment of comparing two homes which are contiguous properties on the same street. Any distortion in assessment ratios arising from neighborhood factors would most plausibly be equivalent for these two homes. Therefore we will describe and remaining difference in assessment ratio as a homeowner effect. We do not observe transactions in sufficient quantity to conduct this analysis using literally adjacent homes. Rather, we will show estimates that appear to converge as we first condition on census tract, and then condition on census block group.

U.S. census tracts are regions of 2,500–8,000 people, with an average of 4,000. Importantly, according to the U.S. Census Geographic Areas Reference Manual, census tracts are initially drawn with the goal of being “as homogeneous as possible with respect to population characteristics, economic status, and living conditions.” This criteria provides additional support for our strategy of attempting to hold neighborhood composition fixed by looking within tract. Table 2 shows the results of a within-tract analysis.³⁴ The homeowner effect for black homeowners is 6.4%, for a black or Hispanic homeowner 5.3%, and for other nonwhite homeowners just under 2%.

As noted, the within-tract analysis seeks to absorb variation in spatial characteristics which drive part of the assessment ratio distortion. However, tracts may be large enough that home prices are not identically affected by local factors. We can conduct the analysis at the block group level. The Census partitions tracts into block groups: regions of 600–3,000 people. This delineation provides an even more defensible setting for our assumption of constant neighborhood characteristics. We obtain block group shapefiles from the U.S. Census and assign all homes in our sample to their corresponding block group (and jurisdiction in keeping with footnote 34.) Table 3 shows the results of a within-block-group analysis. The

dictions. In keeping with our empirical strategy, Figure 10 shows demeaned variation within jurisdiction.

³⁴ As always, our analysis is within jurisdiction. Tracts are sometimes split between jurisdictions. Thus, to be precise, we use jurisdiction-tract-year fixed effects.

estimates are fairly stable relative to the tract-level analysis: the point estimates are 5.9% and 4.85% for black and black or Hispanic homeowners respectively; these are both approximately 50bps lower than the estimates in Table 2. The point estimate for other nonwhite homeowners is almost the same at 1.9%.

We compare columns (1) and (2) in Table 3 with the counterparts in Table 1. For black residents, the homeowner effect is 46% of the overall effect. Considering black or Hispanic residents, the homeowner effect is 49% of the total. For the grouping of homeowners who do not identify as white, black or Hispanic, the homeowner effect is 68% of the total. As we describe in the next section, these homeowners on average face a set of neighborhood characteristics most similar to those faced by non-Hispanic homeowners, and accordingly the neighborhood composition effect is small overall.³⁵

5.2.2 Neighborhood Composition

We next explore the portion of the assessment gap which is conditioned away in the preceding analysis by holding spatial factors constant. Figures 7–10 provide suggestive evidence that racial spatial sorting is relevant for understanding the assessment gap. In each county there is a high tract-level correlation between: (i) highest assessed values relative to market prices and (ii) highest population share of black or Hispanic residents. We establish that this pattern holds in the nationwide data. We estimate the following regression:

$$ar_{i,c,j,t} = \beta brace_{i,c,j,t} + \theta share_{c,j,t} + \epsilon_{i,c,j,t} \quad (6)$$

where ar is the log assessment ratio, i indexes house, j jurisdiction, c census tract, and t year. $Share$ is the tract-level population share for a given racial or ethnic group. Fixed effects are again at the jurisdiction-year level. The results are shown in Table 4. First, the coefficients on demographic shares are all strongly significant, showing that racial composition correlates strongly with the assessment gap. Second, notice that the direct racial/ethnic coefficients in all columns are much reduced relative to our findings in Table 1. In fact, the coefficients are much closer to our estimates of the homeowner effect. This reflects that demographic shares are a strong, though imperfect, statistical proxy for neighborhood factors that correlate with tax-burden variation. We will return to this finding in a number of other ways throughout this section.

A somewhat subtle point is important here. The large estimated effects of demographic

³⁵ Again, this grouping obscures a large amount of underlying racial heterogeneity. We include these results for completeness.

shares in Table 4 are not, by themselves, evidence of meaningful racial inequality. If white and black residents were evenly spatially distributed, the loading on demographic shares in Table 4 would not contribute to average racial disparity.³⁶ Of course it is not the case that homeowner location is randomly assigned. In 2017, the average black resident in the U.S. lived in a tract with 43.5% black share, while the average white resident in the U.S. lived in a tract with 7.2% black share.³⁷ For black or Hispanic residents, the same figures are 56.6% and 17.2% respectively. It is this pattern of residential sorting that, in conjunction with Table 4, implies racial and ethnic inequalities linked to spatial factors.

We are not making any causal claim about the estimates in Table 4. To the contrary, we will now provide evidence that supports the notion that racial and ethnic shares are a statistical proxy for some latent vector of factors which correlate with the assessment gap. Our baseline findings are group-mean differences in the assessment ratio. Any racially-correlated variable which affects the numerator of this ratio (assessments) differentially from the denominator (market prices) will affect the inequality we measure. Our current aim is to show that many neighborhood-level variables generate variation in the assessment ratio. Then, in the next section, we will use hedonic models to be more precise about which category of variables seems to generate the largest mismatch between the two prices.

From the American Community Survey we extract a range of variables observable at the tract-level, and which would plausibly be assumed to affect house prices. We include these variables in equation 6. The results are presented in Table 5. To facilitate interpretation, we scale all variables by their standard deviation, to show the percent change in assessment ratio correlated with a 1 standard-deviation change in a given variable. Again, the equitable taxation null is that all coefficients should be zero. The coefficients on individual homeowner race, which are very similar to Table 4, are shown at the bottom of the regression table.

The surface-level takeaway from this is that many things correlate with the assessment ratio. A positive coefficient represents an increased tax burden, and so Table 5 says not only does higher minority population share correlate with higher tax burden, so does lower median income, higher local unemployment rates, and a larger proportion of residents receiving SNAP benefits.³⁸ Owner percentage and tract level GINI coefficient (a measure of income inequality) are also significantly different from zero. Median age appears to contribute little.

How do we interpret these correlations in the context of the assessment gap? We argue

³⁶ If spatial distribution were truly randomly assigned, then by definition any variation in demographic shares would be statistical noise, and the estimated coefficient should be zero. We are making the point that as residential sorting approaches zero, the inequality implied by any loading on demographic shares also approaches zero.

³⁷ Authors' calculations using American Community Survey data.

³⁸ This is the Supplemental Nutrition Assistance Program, the largest federal nutrition assistance program.

that these patterns arise from market prices being more responsive to neighborhood characteristics than assessed values are. To fix ideas, suppose that assessors impute values as a simple function of home size alone: $A_{i,c,j,t} = f(\text{squarefeet}_i, \#bedrooms_i, \#bathrooms_i)$. It is well established in the housing literature that local amenities are also capitalized into home prices (Roback 1982, Gyourko and Tracy 1991, Cellini et al. 2010). Thus, if $M_{i,c,j,t} = g(\text{squarefeet}_i, \#bedrooms_i, \#bathrooms_i, \text{unemployment}_{c,j,t})$, and the market places a nonzero price on local unemployment, then tract-level variation in unemployment will generate variation in the assessment ratio. Further, if the market hedonic price for unemployment is negative, and if unemployment is correlated with minority demographic share (within jurisdiction), then the mismatch will generate an assessment ratio that is increasing in minority share.

The data is consistent with this very simple framework. We establish this by presenting evidence from two hedonic regressions: one with market values as the dependent variable, and the other with assessed valuations as the dependent variable. Specifically, we specify regressions of the form:

$$y_{i,c,j,t} = \alpha_{j,t} + \beta^y X_{c,j,t} + \theta W_{i,c,j,t} + \epsilon_{i,c,j,t} \quad (7)$$

where $y \in \{A, M\}$, and i indexes home, j government jurisdiction, c census tract, and t year. $X_{c,j,t}$ is a vector of tract-level characteristics, and $W_{i,c,j,t}$ is a (potentially time-varying) vector of home characteristics including square feet, bedrooms, total rooms and flags for various amenities. We are interested in comparing the coefficients on β^A with those on β^M . That is, we are interested in knowing whether hedonic characteristics appear to be *differently* capitalized into market valuations and assessed valuations.

Figure 11 conveys the results of this analysis. Each bar represents the sensitivity of the (log) assessment ratio with respect to a one standard-deviation change of the given variable. At zero, the assessment hedonic model matches the market hedonics. Above (below) zero, the market hedonic prices are larger (smaller) in magnitude than the corresponding assessment hedonic prices. Finally, bars in black are property-level attributes, and bars in red are tract-level attributes. Figure 11 shows that within the context of this hedonic estimation, assessments line up well with market prices on home-level characteristics, but match much less well on neighborhood characteristics. The black bars are all less than 1%: this means that a one standard-deviation shift on any of these dimensions induces less than a 1% shift in the assessment ratio. By contrast, misalignment on tract-level attributes between the assessment and market models is up to an order of magnitude larger. Further, the one variable which receives a greater loading in the assessment model than in the market model

is square feet. Table 6 shows the estimated hedonic prices from both models. Notice that the signs of the coefficients are all relatively intuitive, with the possible exception of owner share. From columns (2) and (4), we can see that assessors clearly do pay attention to neighborhood characteristics in some manner, but don't place *enough* emphasis thereupon. As a whole, the evidence in Figure 11 suggests that assessors: (i) overweight the size of the home, (ii) value other home characteristics fairly precisely, and (iii) underweight local neighborhood composition characteristics.

At a technical level, this underweighting could arise in several ways. All would generate the type of pattern we show here. One possibility is that assessors use hedonic models that include only home attributes and a geographic fixed effect to drive spatial variation in prices. In this case, if the geographic fixed effect is for too broad a region (an entire city, or a quadrant of a city, for example), assessments would be insufficiently high in sub-regions the market values highly and insufficiently low in sub-regions where market prices are low. A similar pattern would result if assessors generate assessments by applying a local growth rate to the prior year's assessment, and that local growth rate is held fixed over a large region (if a single growth rate were picked for an entire city, for example).

While the evidence we provide is consistent with these stories, we are not able to generate a direct empirical test of this hypothesis. Further we cannot disentangle the "mistake" of assessors failing to place enough weight on neighborhood characteristics as the market does, from a deliberate adjustment story where the assessors know how to construct correct valuations but then purposefully distort them in ways that increase the burden on low-income, high-minority, or otherwise economically stressed communities. Empirically speaking, especially at the aggregate level, these would look the same to the econometrician. Either way, the result is a differential in tax burden that arises from neighborhood composition, and creates disparate impact by race.

5.3 Mechanism of Homeowner Effect

The neighborhood composition effect is relatively straightforward. In pure economic terms, this looks like evidence that market prices are more efficient in capitalizing amenities and intangibles than nonmarket (administrative) prices. A race-based differential that attaches to individual homeowners is a little more difficult to explain. A natural intuition might be to think of racially biased assessors. We cannot, in fact, rule this out. However, the practical reality of assessments suggests that assessors are unlikely to know the race of the person within any given home. While it is entirely possible that in some smaller regions the property assessor appears at the front door of the home with a clipboard and a checklist, in larger

regions there are too many properties to make this practical. Automated Valuation Models or Computer Assisted Mass Appraisal is the standard for larger jurisdictions. The International Association of Assessing Officers seems to be the preeminent professional organization in this space, and publishes professional standard guidelines for mass appraisal ([IAAO 2018](#)). The standards essentially outline multivariate regressions using a relatively small vector of property-level characteristics. It is difficult to think of any reason that an assessor performing mass appraisal of numerous properties would know the race of any given homeowner.^{[39](#)}

In every jurisdiction of which we are aware, some process for appealing an assessment exists. In general, this tends to be a bureaucratic process run by some agency of local government.^{[40](#)} A long line of literature in the social sciences suggests a racial component in the extent to which individuals have confidence that public institutions are designed to serve them (extensively surveyed in [Nunnally 2012](#)). This belief may be grounded in reality, or it may be inaccurate but lead to disengagement nonetheless. Therefore, one mechanism we hypothesize and test for is racial differentials in propensity to appeal, likelihood of successful appeals, and degree of reduction conditional on appeal.^{[41](#)} If one group of residents is more effective at reducing assessment growth by navigating the appeals process, this would lead exactly to the wedge between assessments and transacted values that we observe.

We test this appeals channel for the assessment gap in two ways. First, we use a single large county as a case study. We are unaware of any compiled dataset of appeals at a national level. Although the records are quasi-public, appeals records seem to be posted online less often than the tax rolls themselves. Therefore, we obtain a comprehensive record of appeals submitted to the Cook County Assessors Office between 2002 and 2015, courtesy of Robert Ross ([Ross 2017](#)). Covering 1.9M homes and a population of 5.2M (including the city of Chicago), Cook County is the second largest county in the United States. The Cook County records contain the same anonymized property-ID variable as the ATTOM dataset, and therefore is able to be merged directly with our baseline dataset. This yields three

^{[39](#)} The estimates in Tables [2](#) and [3](#) are already holding neighborhood demographics fixed, so any probabilistic inference about race based on local racial demographics is not a plausible explanation for the homeowner effect.

^{[40](#)} Our review of state legal codes suggest that two examples are most common: in one case appeals are made directly to a county assessor's office, and in the other case the state empowers some upstream board of review which has authority to adjust the local assessment.

^{[41](#)} Other scholars have raised this possibility in a property tax setting. Existing work shows a correlation between neighborhood-level demographics and appeal outcomes. [Weber and McMillen \(2010\)](#) also use data from Cook County, covering the period of 2000-2003, along with tract-level demographic data. In a between-tract analysis, they find that high minority share census tracts are correlated with fewer appeal applications and lower success rates. [Doerner and Ihlanfeldt \(2014\)](#) have similar findings in 2005-2009 data from Florida, using a between-block group analysis. To the best of our knowledge, we are the first to use property-level data on homeowner race and ethnicity to conduct a within-neighborhood analysis.

additional pieces of information for each property in Cook County: (i) if an appeal was filed in a given tax-year, (ii) whether the appeal was successful, (iii) if successful, the amount of the reduction.

Cook County has four different channels for appeals: (i) directly through the county assessor's office, (ii) a county board of review, (iii) a state board of review, and (iv) legal appeal through the Illinois circuit court. Staff at the assessor's office tells us that these latter two are most relevant for commercial properties. The data on residential appeals reflects 3.4M total property-level appeals made through the county assessor's office and through the county board of review. Each record tells us if there is a win at the assessor or board of review level, along with the granted reduction. Staff tell us that usually homeowners appeal first to the assessor's office, and then if unhappy with the assessor's decision, may subsequently pursue the appeal at the county board of review. We are unable to distinguish between a homeowner who accepts a first-stage rejection and one who continues but subsequently loses the appeal at the county board stage. As these two venues are tightly grouped within county administration, we will simply denote a "win" as any homeowner who files an appeal and receives a reduction of any amount, regardless of which office approves the reduction.

While Cook County contains many partially overlapping taxing entities, the county is the only body which produces assessments. We are testing the extent to which appeals can explain the 5% portion of the assessment gap driven by the homeowner effect, and thus we will conduct our analysis within block-group-year. We are, therefore, comparing appeal propensity, success, and (conditional) magnitude of reduction between two homeowners from the same block group in the same year. Table 7 shows the results of this analysis. The estimates in columns (1) and (2) use a linear probability model. The specification in column (3) uses the reduction as a proportion of the proposed assessment as the dependent variable. The baseline rate of appeals in Cook County ranges from 10% to 21% annually during this period, with a mean of 14.6%. The estimate in column (1) shows that black homeowners are 84 basis points less likely to appeal. The baseline success rate for assessment appeals in Cook County ranges from 52% to 80% during this period. The mean is 67.4%. The estimate in column (2) shows that black homeowners are 2.2 percentage points less likely to win, conditional on appealing. The mean reduction granted to a successful appeal in this sample is 12.0%. The estimate in column (3) shows that conditional on a successful appeal, black homeowners receive a reduction smaller by .48 percentage points. Results in Table 8 are broadly similar when considering black or Hispanic residents together.

This racial differential in appeals outcomes will, over time, generate different assessment growth rates. White homeowners will appeal with greater frequency and success, which will generate lower assessment growth relative to black or Hispanic homeowners. Absent

other data on appeals, we cannot directly test the assessment appeals channel in other jurisdictions. We can, however, test whether the national data shows evidence of the patterns which this channel would generate. We can exploit the time-series structure of assessments in the ATTOM dataset to ascertain whether assessment growth varies by homeowner race or ethnicity. Due to our focus on the assessment ratio, all baseline findings consider homes where we observe an assessment and a market transaction within the same period (year). Assessments are produced annually however, regardless of whether a transaction occurs.⁴² Thus we can test for a racial differential in the trajectory of assessments over time.

We will exploit the fact that for a large number of homes in our sample, the racial ownership changes pursuant to a transaction. This permits us to estimate a generalized difference in differences model:

$$y_{i,c,j,t} = \alpha_i + \gamma_{c,j,t} + \beta^r race_{i,c,j,t} + \epsilon_{i,c,j,t} \quad (8)$$

Each property in this sample is sold at some point. β^r is identified from properties which undergo a change in racial ownership as a consequence of the transaction. Property level fixed effects absorb the between-home variation, and jurisdiction-tract-year fixed effects absorb local housing market variation. The identifying assumption is that within year and census tract, homebuyer selection into properties is orthogonal to future home price shocks. With this assumption, β is the causal effect of racial ownership on assessments.

Table 9 shows the results. As homeowners typically can appeal their assessments each year (or as frequently as new assessments are generated), the channel we posit is most relevant to growth. Accordingly columns (1) and (2) use the assessment growth (log differences) as the dependent variable. The coefficient in column (1) says that assessment growth is 7bps higher when a black person owns a property, relative to when a white person owns the same property. This is significant only at the 10% level. For black or Hispanic residents the difference in growth is 41bps, and is strongly statistically significant. Given that our sample spans only 13 years, and that an initial transaction is necessary to pin down the race and ethnicity of the homeowner, estimating growth rates may be straining the data even though the sample is very large. In columns (3) and (4) we use (log) levels as the dependent variable instead. The level difference is 29bps and 79bps respectively. This is consistent with the growth evidence. Within property, assessment levels are higher for minority residents.

⁴² To the best of our knowledge, property taxes are paid annually in every jurisdiction. Thus, for purposes of producing a bill, there is an assessment for every tax year; this is what we observe in the ATTOM dataset. In jurisdictions that revise assessments less often than annually, the assessment remains static for some number of years (typically 1–2 years, but sometimes longer).

Given the length of our sample, the estimates in columns (3) and (4) should be thought of as reflecting two to three assessment cycles, which suggests reasonable consistency between the growth estimates and level estimates.

Our growth estimates may be somewhat small relative to the magnitude of the homeowner component of the assessment gap, which was on the order of 5–6%. The more precise (and larger) estimates of column (2) would suggest that the 5% effect (column 2 of Table 3) would be generated in 11.5 years. This is a very reasonable figure for median homeowner tenure.⁴³ The growth differential estimated in column (1), however, would require 80 years to generate the corresponding homeowner effect for black residents. This may suggest that appeals are not the only mechanism in play.

To the extent that assessors do not, in fact, know the race of the homeowner – which is more likely to be the case than not – we argue that Table 9 provides strong indirect support for an appeals channel. It is difficult to think of another plausible driver. Any other explanation would require ex-ante racial sorting on *future* assessment growth. The alternate hypothesis would be in effect that white homebuyers are more likely to select properties that will face a negative assessment shock in the future. We find this less likely than an assessments channel, but acknowledge that the issue remains open for future research.

5.4 Heterogeneity by Racial Attitudes and Demographic Composition

It is natural to wonder how the assessment gap relates to racial attitudes. For both channels outlined above, active expression of bias is not necessary, but neither can we rule it out. We use two measures of racial animus developed in [Stephens-Davidowitz \(2014\)](#) to split our sample into regions of high and low racial prejudice. In each sub-sample, we estimate the overall assessment gap and the homeowner effect. The racial animus measures are derived from the regional intensity of Google searches that contain the most offensive epithet for African-American people. One measure is produced at the state-level, and the other at the media-market level. For the latter, we use a Nielsen crosswalk to assign the media market measure to counties. We then split our sample along the median of each measure, and estimate the assessment gap in a pooled regression. As the measure is designed to capture prejudice towards African-Americans, we estimate the assessment gap only for black homeowners, and not for other groupings of minority residents.

Table 10 shows the results. Using either measure, the assessment gap is significantly

⁴³ The ACS data implies a median tenure of approximately 12 years.

larger in high-animus regions. This holds both in the overall estimates shown in columns (2) and (4), and in the homeowner effect estimates in columns (3) and (5). In regions of below-median prejudice, the assessment gap is still economically and statistically significant. Several mechanisms could lead the assessment gap to be increasing in racial animus. In higher animus regions, minority residents may be marginally less likely to appeal property assessments, or less likely to succeed in that appeal. Or, high animus regions may lead to increased racial residential sorting and a larger market-price capitalization of racially correlated factors, which would also lead to the pattern observed here.

We also split the national sample into quintiles based on minority population share at the county-level. The first quintile contains counties with the smallest minority share, and the 5th quintile is comprised of counties with the largest. We estimate the assessment gap in each of these sub-samples. Figure 12 shows results from these regressions graphically, and Table 11 shows the regression estimates. The assessment gap is clearly increasing in minority population share. Since we have shown that a large portion of the assessment gap is linked to spatial sorting, this finding is unsurprising: it has been documented that spatial sorting increases as minority population increases (Card et al. 2008).

5.5 Effective Tax Burden

Our last set of results link the assessment gap distortion with actual higher taxation. As a matter of theory, any wedge between assessments and market prices must create a distortion in an ad valorem tax. We are able to observe taxes paid, and therefore can provide the empirical evidence showing that this theoretical relationship does, in fact, hold. Thus far, our focus on assessment ratios has been very deliberate. Assessed values and market prices are observable by the econometrician with little ambiguity. Taxes are more complicated, chiefly due to exemptions.

Every state provides for a variety of tax exemptions in state legislative codes. Most localities have further autonomy to create exemptions. A common example would be a principal residence exemption: Michigan, for example, exempts primary homes from school taxation up to the amount of 18 mills (180bps).⁴⁴ Another very common exemption holds for residents of retirement age: New York State permits an exemption of up to 50% for residents over 65 whose income is between \$3,000 and \$29,000.⁴⁵ Within these parameters, local units have autonomy to select the precise cutpoints. While these are relatively straightforward, many exemptions are much more complicated. Even at a state level, the list of exemptions tends

⁴⁴ Michigan Compiled Laws, Section 211.7cc and 380.1211.

⁴⁵ <https://www.tax.ny.gov/pit/property/exemption/seniorexempt.htm>.

to be very long and complex. With tens of thousands of local authorities also potentially creating additional exemptions, even observing these exemptions becomes a significant challenge. While the ATTOM data includes a field for exemptions, it is unclear how consistently or accurately this data is reported. We show results: (i) using the reported tax bill directly, and (ii) adding back in the reported exemptions to create a pre-exemption tax bill.

Exemptions matter in general because spatial distribution of the exemptions may very well be correlated with racial demographics. If some parts of Florida have more elderly white residents than young black residents, the exemption policy itself will create something that looks like a distortion in the tax burden, but which is entirely consistent with the legislative intent and public administration of the tax system. We are unable to observe, and thus control for, age of the homeowner – let alone any other individual-level drivers of more complicated exemption policies. The strength of considering the assessment ratio is that none of these confounding factors matter. Using tax dollars paid, we are less able to rigorously strip out potential confounding factors.

Another complicating factor is partial-year tax bills. In some jurisdictions the homeowner of record on a certain date is liable for a full year’s worth of property taxes. In others, a partial year of ownership would result in a tax bill spanning only that portion of the year. We do not observe this policy choice at a local level. As we need market prices to determine an effective tax burden, we have another source of bias if race correlates with any propensity to sell in any given year. To provide robustness around this issue, we will compute effective tax burden during the sale year, as well as one year before and one year after sale.

We first estimate the pass-through of the assessment ratio to the effective tax rate. We regress the log effective tax rate on the log assessment ratio. The mechanics of property tax administration would suggest a coefficient of 100%, unless homeowners have not fully exhausted available exemptions. If a region permits homeowners to deduct \$5,000 from the assessed value of their primary residence before computing the tax bill and many homes are assessed at less than \$5,000 then the pass-through would be less than 100%. Table 12 shows these estimated pass-through rates. Column (1) presents estimates for all homeowners in aggregate, and columns (2) and (3) show results by racial and ethnic grouping. Results for black residents alone are very similar, and we do not include them here. Columns (1) and (2) use the actual tax bill. The pass-through is 99%, which closely matches the prediction, and suggests that some homeowners are perhaps inframarginal with respect to deductions. Column (3) uses the computed pre-exemption tax bill. Here the estimates are lower. If anything, theory would suggest this number should be closer to 100% than post-exemption figures. We see little reason for the lowered estimates. We take this as an additional reason to be wary of the reported exemption data. Across columns (2) and (3), differences by racial

or ethnic identity are not evident.

Table 13 directly estimates racial differentials in effective tax rate during the sale year. For black residents, we estimate an effective tax rate that is 14.9% higher in the actual tax bill, and 12.2% higher with exemptions added back. This closely brackets the 12.6% excess burden suggested by the assessment gap. Considering black or Hispanic residents together, we find a 11.4% higher effective tax rate from tax bills, and 8% increase with observed exemptions added back. Again, this brackets our assessment gap estimate of 9.8%. For the grouping of other nonwhite homeowners, the estimates in columns (5) and (6) again align tightly to our baseline assessment gap estimate. Appendix Tables A5 and A6 show very similar patterns using tax bills one year on either side of the sale.

5.6 Additional Discussion: The Role of Market Prices

Market prices are central to our empirical strategy. As discussed in sections 2 and 3, this reflects the intention of the property tax as outlined by authorizing legislation. Accordingly, we interpret realized market prices in an arms-length transaction as the appropriate basis for taxation. But what if market prices are “wrong”? A racial differential in transacted prices would also generate a wedge between assessed values and market values, even if assessments perfectly reflect true latent value. It is not immediately obvious how to think about inequality generated by such a mechanism. On the one hand, taxation in general is usually applied to realized financial flows rather than some latent value. On the other hand, it hard to imagine tasking assessors with incorporating homeowner-specific factors that affect market prices. In this section, we provide some evidence on the potential magnitude and direction of inequality driven by racial or ethnic differences in transacted prices.

Many economics papers have explored this possibility. Bayer et al. (2017) uses very similar housing microdata to the ATTOM dataset used in this paper, and finds that black and Hispanic *buyers* pay a premium of around 2%. This effect is positive across virtually all racial and ethnic combinations of buyers and sellers, and is largest for within-race transactions (black seller and black buyer; or Hispanic seller and Hispanic buyer). In US housing markets, the majority of transactions occur within-race. Therefore the Bayer et al. (2017) finding would suggest that minority assessment ratios in our sample (which are associated with the race and ethnicity of the home *seller*) may be understated by 2%. This would imply that racial or ethnic differences in transacted prices *lower* our estimates of inequality.

Bayer et al. (2017) uses a within-property analysis and restricts attention to four large metropolitan areas to obtain sufficient transaction density. One embedded assumption is that home characteristics stay constant (and are therefore absorbed by the property-level

fixed effect). We add additional evidence using a slightly different methodology. While the ATTOM dataset does provide time-varying home characteristics, only macro-level attributes are captured: number of bedrooms, square feet, etc. Assessors typically track major home improvements closely. However, the trigger for updating assessor data is generally a construction permit filed with some local public bureaucracy. We therefore would not observe an indicator for, say, replacing kitchen floor tiles, or some other relatively minor improvement which nonetheless would likely impact market price. To address this, we test for racial and ethnic differences in transaction prices which are not predicted by local housing market conditions.

In the set of homes which sell more than once, we define P_0 as the first transaction price. We then form a predicted selling price:

$$\hat{P}_{i,t} = P_{i,0} * \frac{HPI_{z,t}}{HPI_{z,0}} \quad (9)$$

where $HPI_{z,t}$ is a zip-code level home price index for time t.⁴⁶ We then run the following regression:

$$\ln(P_{i,t}) - \ln(\hat{P}_{i,t}) = \gamma_{bg,t} + \beta^r \text{seller race}_i + \epsilon_{i,z,t} \quad (10)$$

where γ_{bg} is a census block group fixed effect. The left hand side is the difference between realized transacted prices and predicted transaction prices. We include a fixed effect at the block group level to account for spatial errors generated by use of a zip-code HPI.⁴⁷. The coefficients on our categorical *seller race* variable are estimates of racial and ethnic differences in transacted prices which are not explained by local housing market conditions.

Table 14 shows the results. We estimate that black sellers receive 2.2% more than white sellers within the same census block group. Considering black or Hispanic sellers together, the estimated premium is 3.3%. This evidence lines up very closely with the results presented in [Bayer et al. \(2017\)](#). The difference in transacted prices could arise from differential propensity to improve or maintain property, or from a range of housing market frictions. No matter the reason, these results suggest to the extent that a racial differential in market prices exists, realized market prices are slightly higher for minority sellers. This would lead to a *lower* assessment ratio for minority sellers, which means that our estimates of inequality are, if anything, biased downwards on the order of 2-3%.

⁴⁶ This use of zip-code level home price indexes holds much in common with the policy approach discussed at greater length in section 6. As in that section, we obtain zip-code HPI measures from Zillow.

⁴⁷ As discussed further in Section 6, zip codes are relatively large regions.

5.7 Additional Discussion: Race vs Income/Wealth

Our baseline estimates of inequality all condition on taxing jurisdiction (annually), but do not include any other control variables. This is intentional. Taxing jurisdictions are formed to create regions where every resident faces the same level of intended taxation. Our equitable taxation null shows that any within-jurisdiction difference in assessment ratios represents an inequality in tax burden. As noted, this is a concept of inequality that aligns tightly to the legal standard of disparate impact. Thus, in this setting, the unconditional difference (within jurisdiction) is the primary statistic of interest.

However, as we have noted, racial and ethnic wealth disparities are among the most persistent and salient stylized facts in household finance. This begs the question of whether the assessment gap simply arises because the U.S. property tax is more regressive than previously understood. This would imply that differences in assessment ratio relate only to income, and that the assessment gap simply appears racially tinged when viewed through the lens of race and ethnicity. We believe the data strongly rejects this notion.

First, our estimates of the homeowner effect are within census block group. This spatial conditioning is a fairly robust non-parametric control for income in many parts of the U.S. Further, our estimates of a racial difference in assessment growth rates are within property. To the extent that choice of home value is also a statistical proxy for income or wealth, this further suggests that the homeowner effect is weakly linked to income.

Our findings on neighborhood composition show that many highly local features are under-capitalized into assessed valuations, thereby generating inequality in tax burden. Tables 5 and 6 suggest that neighborhood income is one of these features – at least in a linear specification. So part of our findings do, in fact, suggest that a portion of the inequality we document relates to regressive features of the property tax. However, as table 5 shows, not only are individual racial and ethnic covariates still large once neighborhood traits are added, minority share is also a highly significant predictor of the assessment gap even after controlling for a range of socioeconomic factors (again, linearly).

We add another piece of suggestive evidence by estimating the assessment gap within ventile of median tract-level income. If the assessment gap were generated primarily by income-related factors, then we should find little evidence thereof within groups equalized by income or wealth. Figure 13 shows the estimated assessment gap for black homeowners within income quantile. Two features of this figure are most salient. First, the estimated assessment gap is economically significant within all quantiles. From approximately the median tract through the highest income tract, the level of inequality is fairly stable and is approximately 5%, which matches the magnitude of the homeowner effect. The second salient feature is that the assessment gap is sharply increasing in lower income quantiles. This shows

that inequality generated by the assessment gap is heterogenous in income. However, this feature also strongly rebuts the notion that the assessment gap arises primarily from income-related factors. The lowest quantiles in Figure 13 are essentially comparing (jurisdictionally demeaned) assessment ratios for the poorest white residents with the poorest black residents. If income were the primary mechanism, these estimates should be close to zero. That they are starkly increasing shows that property assessments are much more misaligned to market values for low-income minority residents than they are for low-income white residents. Figure 14 shows the corresponding analysis for black or Hispanic homeowners.

6 Policy Corrections

In this section, we discuss a potential approach to address the assessment gap. The inequality we document stems from a wedge between market prices and assessments. Having carefully documented the extent and magnitude of the distortion, it is natural to ask how easily the problem could be fixed. Perhaps it is the case that market prices are so sensitive to geographic variation, and property prices so temporally unpredictable that even the most-skilled and attentive assessors office would not be able to equalize tax burdens by racial status. In this section we show that a relatively simple approach can address a large portion of this inequality.

As more than half of the assessment gap relates to mispricing of local characteristics, we explore whether small-geography home price indexes (HPIs) can be used to reduce inequality. We use zip-code level HPIs to produce imputed assessments, and then compare the racial variation in assessment ratios constructed with our constructed assessments to the variation using true assessments. We find this simple procedure reduces inequality by 55–70%. The average zip code is about twice as large as a census tract. We conjecture that more geographically precise HPIs would be even more effective in removing assessment ratio variation.

We use publicly available zip-code level HPIs from Zillow to construct assessments. Zillow constructs these HPIs monthly for 15,500 zip codes. This covers 84% of the U.S. population.⁴⁸ As some transaction density is needed for a sample size sufficient to produce a reasonable HPI index, these zip codes are highly skewed towards more populous urban areas. The monthly time-series from 1996 can be directly downloaded from Zillow’s website at no cost. Zillow began providing these indexes in 2006, and has backwards constructed them to 1996. Zillow has also been increasing its coverage over time.

⁴⁸ Author’s calculations using 2010 decennial census data.

We construct synthetic assessments using the zip-code HPIs. The algorithm for a synthetic assessment is simple: in any zip code, we take the first observed transaction price and allow this to be the assessment in the month-year of sale. Then we grow that assessment according to the relevant monthly HPI. That is:

$$\hat{A}_{i,j,z,t} = M_{i,j,z,0} \frac{HPI_{i,j,z,t}}{HPI_{i,j,z,0}} \quad (11)$$

where 0 denotes the base month-year of the 1st transaction, z denotes zip code, and $M_{i,j,z,0}$ is the observed transaction price in the base year.

We next test the inequality which would be generated by using these synthetic assessments as the basis for property taxation. To do this, we apply the algorithm to carry the synthetic assessment forward in time until we arrive at the month-year of a subsequent transaction. We then form a synthetic assessment ratio at that time t by taking the log difference between our synthetic assessment and the observed transacted price: $\hat{ar}_{i,j,z,t} = \log(\hat{A}_{i,j,z,t}) - \log(M_{i,j,z,t})$. We evaluate the success of this algorithm for generating assessments by comparing inequality in synthetic assessment ratios to inequality in the realized assessment ratios. Because this simple approach requires two transactions, and is by construction limited to the zip codes that Zillow covers, we end up with a significantly smaller subsample of 2.1M homes. We first document that the assessment gap still exists – and looks similar – in this subsample. Then we document that using synthetic assessments reduces inequality by 55–70%.

The first three columns of Table 15 show the assessment gap in the subsample covered by Zillow HPIs. Magnitudes are similar to our baseline findings. The figures in columns (1) and (2) are respectively 1.7% and 1.4% larger than the findings in Table 1. For minority homeowners who neither identify as black nor as Hispanic, the estimated effect is nearly the same. Columns (4)–(6) repeat the same regressions using our synthetic assessments. A perfect procedure would produce zeros on the racial and ethnic variables. The synthetic assessments completely reverse the assessment gap, and in fact overshoot. The estimates in columns (4)–(6) of Table 15 reflect a *lower* tax burden on minority residents. Of course this is also an inequality in the tax burden. However, the overall distortion is much smaller in magnitude: 4.1% for black homeowners, 5.1% for black or Hispanic homeowners, and effectively zero for other nonwhite homeowners.

Two things are worth emphasizing here. One is that such a straightforward approach is only feasible if some valid HPI exists for small geographic regions. We use Zillow’s zip-code HPIs to demonstrate that inequality can be reduced by using publicly available, easy to obtain data. Zip codes are, however, well known to be formed with little consideration for

the institutions and characteristics of the underlying geography. Also, the average zip code contains 9,000 people. This is relatively large: our results suggest that there is meaningful spatial variation between tracts, which are less than half this size on average. We think this is likely to be one important reason that this simple implementation still generates a 4-5% racial difference in assessment ratios. The discussion in section 5.6 also suggests that a racial or ethnic difference in transaction prices could explain 2-3 percentage points of the remaining inequality. In addition, as a practical matter, assessment values need to be set at the beginning of the tax year, and sales may occur at any time during the next 12 months. Accordingly, racial sorting into areas of higher or lower growth would cause some amount of measured inequality in the realized assessment ratio to arise within the year. To see how important this channel would be, we reproduce a set of synthetic assessments where the assessment is set annually in January of each year. Every transaction then includes up to 12 months of home price growth which is not reflected in the assessment. Appendix Table A7 shows results from this exercise. The estimates are almost unchanged.

The second point of emphasis is that our procedure uses an observed transaction price for the base year value. In order to apply to all properties within a jurisdiction, assessors would need some method for imputing price for properties which have not sold during the period spanned by the HPI. Our neighborhood composition findings suggest that this too, will require assessors to permit prices to vary between small geographic regions. Racial equity in the initial values is empirically observable and testable. So assessors should be able to iterate a model for initial pricing to land on an equitable distribution of base-year assessments, and then grow those by using some HPI index.⁴⁹ The point remains that assessors can make significant strides towards equity by linking assessment growth to small geographic regions within their jurisdiction.

7 Conclusion

We document widespread racial inequalities in the U.S. property tax burden. The residential property tax is intended to be an ad valorem levy on the fair market value of the owned asset, yet tax bills are generated as a function of a policy rate and an assessed value. Thus, any wedge between assessed valuations and market prices creates some deviation away from a fair tax benchmark. While local jurisdictions have free choice of a scaling parameter for assessments relative to market prices, the realized ratio within jurisdiction should always be

⁴⁹ This is, in fact, not particularly dissimilar from the process advocated by IAOO (2018) and other professional guides. However the bulk of this paper serves to show that regardless of process, the outcomes articulated in standards like these are not being widely achieved.

constant across properties. We obtain shapefiles for a comprehensive set of local governments along with other quasi-governmental entities that levy taxes. We associate each home with all governments that contain it. Within these taxing “jurisdictions,” defined as a unique set of overlapping governments, the assessment ratio should be constant. Our first major finding is to document a nationwide assessment gap: assessment ratios are on average higher for minority homeowners. Holding jurisdiction – and thereby public services, intended taxation and local assessment practices – fixed, the average assessment gap for a minority resident in our sample is 10–13%.

We decompose this finding into two components. We show that approximately half of the overall effect, 5–6%, remains even within very small geography. We hypothesize that the main channel for this effect is racial differentials in property tax appeals. We use administrative data from Cook County, the second largest county in the US, to demonstrate that such racial differentials can exist: in Cook County, minority residents are 1% less likely to appeal; are 2% less likely to win an appeal; and conditional on success, receive a 2–3% smaller reduction. We then exploit racial changes in ownership around property transactions to test for racial differentials in national assessment trajectories, and find patterns consistent with an appeals mechanism.

We show that the remaining half of the assessment gap can be explained by between-neighborhood variation. Residential sorting by race in the U.S. means that the average black or Hispanic resident faces a different set of local attributes than a white resident does. Market prices appear to be substantially more sensitive to a wide range of observable neighborhood characteristics than assessed valuations. We use hedonic regressions to show that market prices and assessed values align well on home-level attributes, but diverge on tract-level characteristics. This mismatch, along with residential segregation patterns, generates the other 6–7% of the total tax burden inequality.

Last, we demonstrate that these distortions can be fixed by a relatively simple procedure. Our results suggest that it is important for assessors to recognize that market prices are highly sensitive to local conditions, in ways that correlate with race. Accordingly, assessed valuations should reflect price dynamics at a narrow geographical level. We obtain zip-code-level home price indexes. We use these to produce synthetic assessments in the simplest way possible: when a transaction occurs, that becomes the assessment, and from there it evolves in direct proportion with the monthly zip-code HPI. Using a subsample of homes for which we observe two transactions, we create the synthetic assessments at the first sale, and model them forward to the second sale. At the second sale, we form a synthetic assessment ratio, and use these ratios to test for an assessment gap. The simple synthetic approach reduces the overall inequality by 55–70%, and in fact, flips incidence: the remaining inequality is a

lower tax burden on minority residents.

Our baseline findings establish that minority residents in the U.S. face a higher property tax burden than their nonminority neighbors. Although the professional standards for the appraisal industry emphasize horizontal and vertical equity of taxation, the reality of property tax administration in the U.S. is that more jurisdictions fail to achieve this equity than not. Increased taxation clearly represents, in the most literal sense, an incremental cost faced by minority families and an additional impediment to minority wealth building. We know already that there are very striking racial wealth disparities in the U.S. – especially between black and white residents. The inequality we document in taxation is a direct, ongoing, and current source of fiscal headwinds for minority families. We estimate an additional burden of \$300–\$390 per year for the median black or Hispanic family, and up to \$790 for families affected at the 90th percentile of the assessment gap. Residents of local governments implicitly enter a contract agreeing to a given level of taxation in exchange for a bundle of public amenities. Nearly every homeowner in the U.S. faces a property tax, and this large-scale shifting of tax burden onto minority residents violates the notions of equity embedded in the implicit contracts that residents make with local governments.

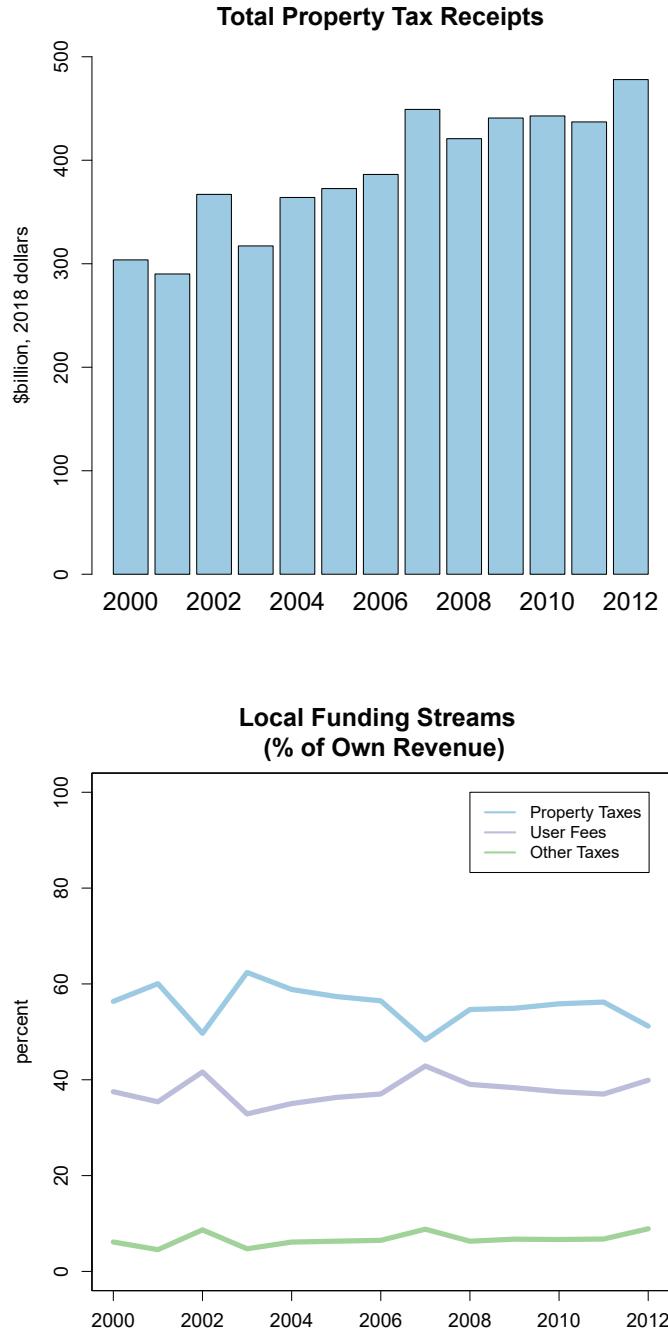
References

- Ananat, E. O. (2011). The wrong side(s) of the tracks: The causal effects of racial segregation on urban poverty and inequality. *American Economic Journal: Applied Economics*, 3(2):34–66.
- Arnold, D., Dobbie, W., and Yang, C. S. (2018). Racial bias in bail decisions. *The Quarterly Journal of Economics*, 133(4):1885–1932.
- Atuahene, B. (2017). Our taxes are too damn high: Institutional racism, property tax assessments, and the Fair Housing Act. *Nw. UL Rev.*, 112:1501.
- Atuahene, B. and Berry, C. (2019). Taxed out: Illegal property tax assessments and the epidemic of tax foreclosures in Detroit. *UC Irvine Law Review*, 9(4):847.
- Bartlett, R., Morse, A., Stanton, R., and Wallace, N. (2018). Consumer-lending discrimination in the era of FinTech. *Unpublished working paper. University of California, Berkeley*.
- Bayer, P., Casey, M., Ferreira, F., and McMillan, R. (2017). Racial and ethnic price differentials in the housing market. *Journal of Urban Economics*, 102:91–105.
- Bayer, P., Ferreira, F., and McMillan, R. (2007). A unified framework for measuring preferences for schools and neighborhoods. *Journal of political economy*, 115(4):588–638.
- Bayer, P. and McMillan, R. (2005). Racial sorting and neighborhood quality. Technical report, National Bureau of Economic Research.
- Bertrand, M. and Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013.
- Card, D. and Lemieux, T. (1996). Wage dispersion, returns to skill, and black-white wage differentials. *Journal of Econometrics*, 74(2):319–361.
- Card, D., Mas, A., and Rothstein, J. (2008). Tipping and the dynamics of segregation. *The Quarterly Journal of Economics*, 123(1):177–218.
- Card, D. and Rothstein, J. (2007). Racial segregation and the black–white test score gap. *Journal of Public Economics*, 91(11-12):2158–2184.
- Cellini, S. R., Ferreira, F., and Rothstein, J. (2010). The value of school facility investments: Evidence from a dynamic regression discontinuity design. *The Quarterly Journal of Economics*, 125(1):215–261.
- Charles, K. K. and Guryan, J. (2008). Prejudice and wages: An empirical assessment of Becker’s The Economics of Discrimination. *Journal of political economy*, 116(5):773–809.

- Charles, K. K. and Hurst, E. (2002). The transition to home ownership and the black-white wealth gap. *Review of Economics and Statistics*, 84(2):281–297.
- Chetty, R., Hendren, N., Jones, M. R., and Porter, S. R. (2018). Race and economic opportunity in the United States: An intergenerational perspective. Technical report, National Bureau of Economic Research.
- Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *The Quarterly Journal of Economics*, 129(4):1553–1623.
- Cutler, D. M., Glaeser, E. L., and Vigdor, J. L. (1999). The rise and decline of the American ghetto. *Journal of political economy*, 107(3):455–506.
- Doerner, W. M. and Ihlanfeldt, K. R. (2014). An empirical analysis of the property tax appeals process. *Journal of Property Tax Assessment & Administration*, 11(4):5–34.
- Donohue III, J. J. and Heckman, J. (1991). Continuous versus episodic change: The impact of civil rights policy on the economic status of blacks. Technical report, National Bureau of Economic Research.
- Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2018). Predictably unequal? The effects of machine learning on credit markets. *The Effects of Machine Learning on Credit Markets (November 6, 2018)*.
- Gittleman, M. and Wolff, E. N. (2004). Racial differences in patterns of wealth accumulation. *Journal of Human Resources*, 39(1):193–227.
- Gyourko, J. and Tracy, J. (1991). The structure of local public finance and the quality of life. *Journal of Political Economy*, 99(4):774–806.
- IAAO (2018). Standard on automated valuation models. *International Association of Assessing Officers*.
- Kahrl, A. W. (2016). The power to destroy: Discriminatory property assessments and the struggle for tax justice in Mississippi. *Journal of Southern History*, 82(3):579–616.
- Kleinberg, J., Ludwig, J., Mullainathan, S., and Sunstein, C. R. (2018). Discrimination in the age of algorithms. *Journal of Legal Analysis*, 10.
- Knowles, J., Persico, N., and Todd, P. (2001). Racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy*, 109(1):203–229.

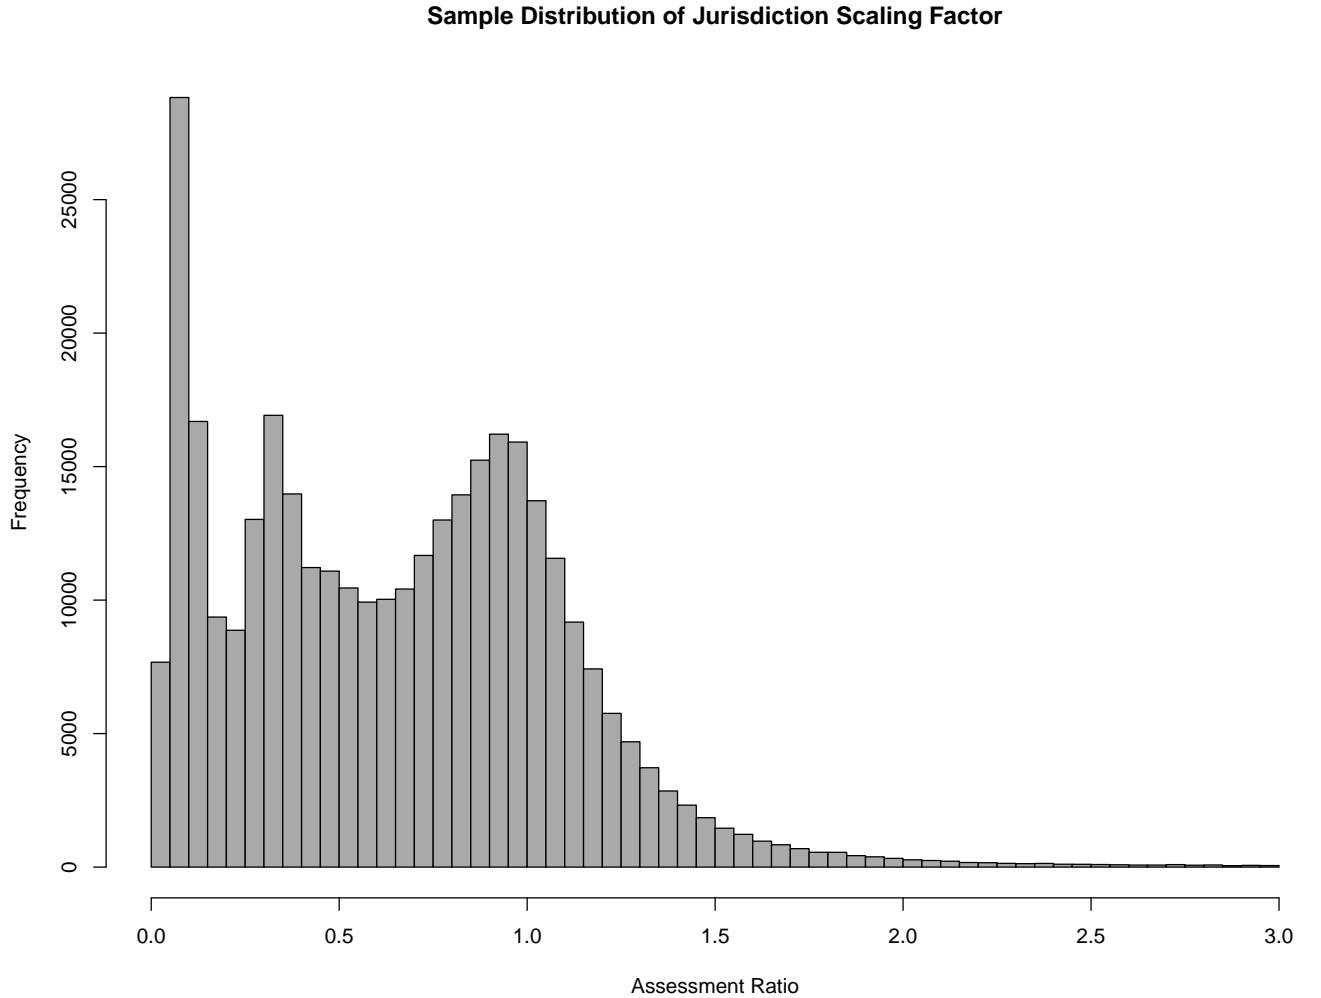
- Lichter, D. T., Parisi, D., Grice, S. M., and Taquino, M. C. (2007). National estimates of racial segregation in rural and small-town America. *Demography*, 44(3):563–581.
- Logan, T. D. and Parman, J. M. (2017). The national rise in residential segregation. *The Journal of Economic History*, 77(1):127–170.
- Massey, D. S. and Denton, N. A. (1993). *American apartheid: Segregation and the making of the underclass*. Harvard University Press.
- Nunnally, S. C. (2012). *Trust in Black America: Race, Discrimination, and Politics*. NYU Press.
- Roback, J. (1982). Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6):1257–1278.
- Ross, R. (2017). The impact of property tax appeals on vertical equity in Cook County, IL. *University of Chicago, Harris School of Public Policy Working Paper*.
- Rothstein, R. (2017). *The Color of Law: A Forgotten History of How Our Government Segregated America*. Liveright Publishing.
- Schulman, K. A., Berlin, J. A., Harless, W., Kerner, J. F., Sistrunk, S., Gersh, B. J., Dube, R., Taleghani, C. K., Burke, J. E., Williams, S., et al. (1999). The effect of race and sex on physicians' recommendations for cardiac catheterization. *New England Journal of Medicine*, 340(8):618–626.
- Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, 118:26–40.
- Weber, R. N. and McMillen, D. P. (2010). Ask and ye shall receive? predicting the successful appeal of property tax assessments. *Public Finance Review*, 38(1):74–101.
- Williams, D. R. (2012). Miles to go before we sleep: Racial inequities in health. *Journal of Health and Social Behavior*, 53(3):279–295.

Figure 1: Sources of Funding for Local Governments



Note: This figure shows total property tax receipts (top) for local units of government, and a breakdown of revenue sources (bottom) for 2000 to 2012. Data is from Census of Governments. Full census years are 2002, 2007, and 2012. In all other years, only larger local governments are surveyed. This subset still represents 80-90% of total government budgets, but the omission of smaller governments causes mechanical spikes and dips in 2002, 2007 and 2012 in both figures. The bottom graph shows the average composition of local “own revenue,” which is the portion of the budget that the local government can directly affect by policy choices. This excludes intergovernmental transfers from state and federal levels of government. Property tax figures shown are for all property taxes, residential and commercial.

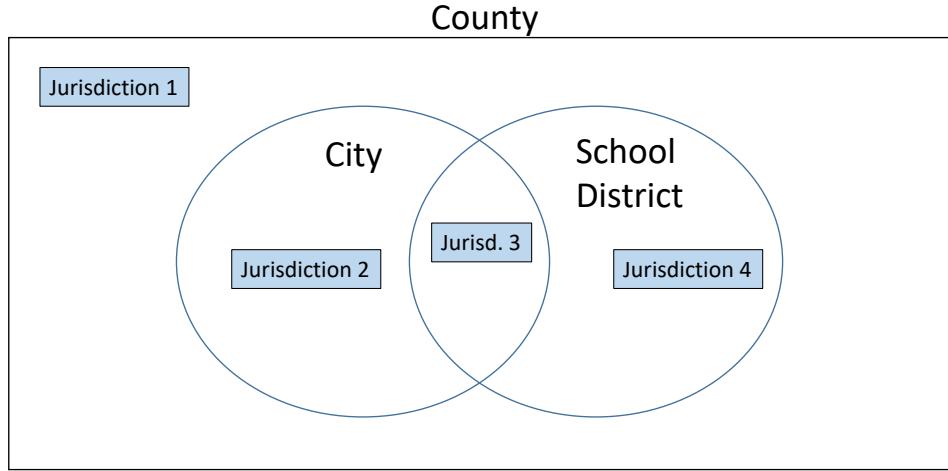
Figure 2: Sample Distribution of Local Scaling Factor



Note: This figure shows the mean *realized* jurisdiction assessment ratio by jurisdiction-year. If assessors make no *average* mistake jurisdiction-wide, then this realized assessment ratio will be equal to the intended scaling factor. In our setting, inequality is a relative difference in assessment ratios between groups within jurisdiction. Therefore, deviations from realized mean are the relevant statistic. If an entire jurisdiction targets a 40% assessment ratio, but realizes a 50% assessment ratio for everyone, this will affect all residents proportionately. Such an outcome may have implications for total revenue raised, but does not represent a source of inequality within jurisdiction. Also note that a jurisdiction-wide wedge between intended and realized scaling factor may not have implications for revenue raised: in many locations the *amount* of intended spending is the politically established choice, and then the aggregate budget is simply divided by aggregate assessed property values to generate a policy tax rate that will be applied to each individual assessment. In this case, if average assessments are higher (lower) across the entire jurisdiction, then implied property tax rates will mechanically be lower (higher) in an exact offset.

Figure 3: Taxing Jurisdiction Stylized Examples

Panel A

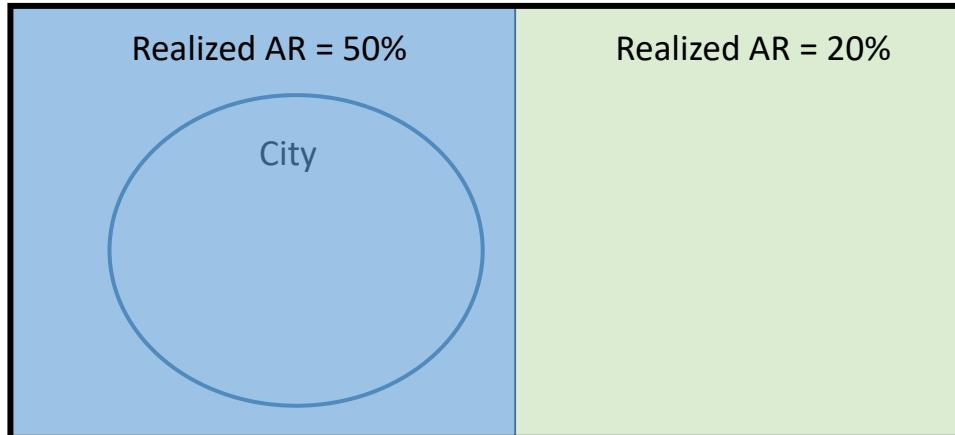


"Jurisdiction":

Region touched by a unique network of overlapping governments

Panel B

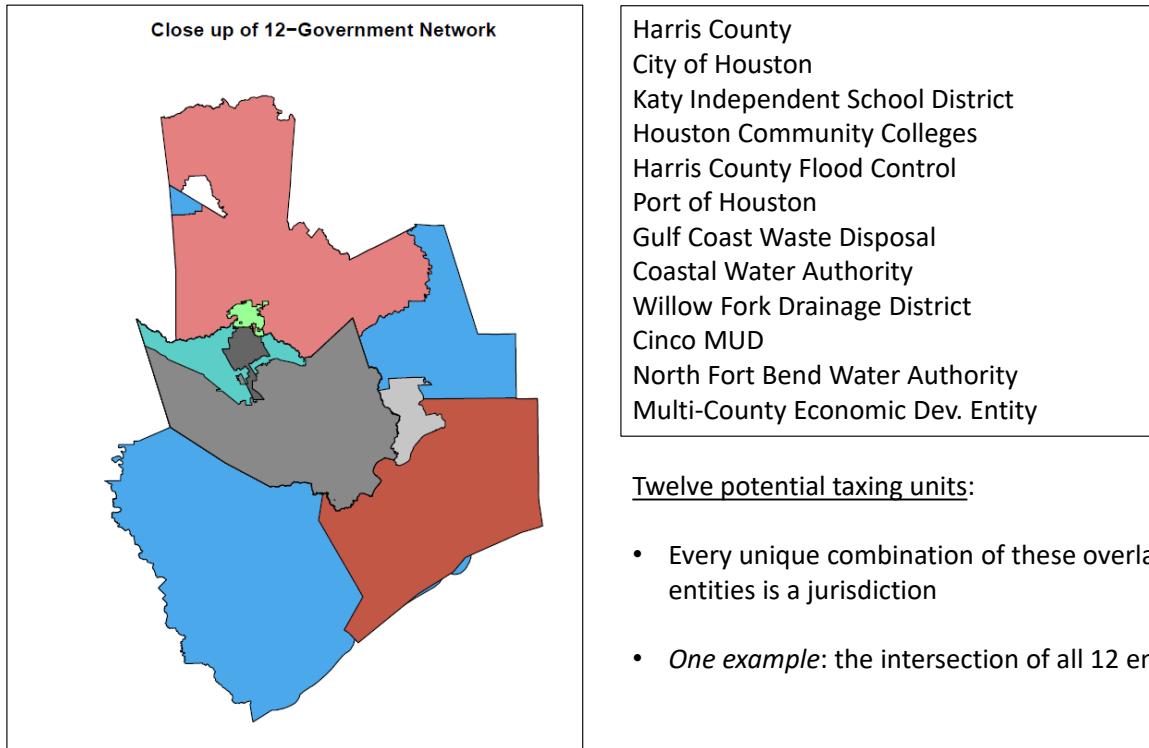
County: Target AR 40%



- 1) Inequality in county tax
- 2) But no inequality in city tax

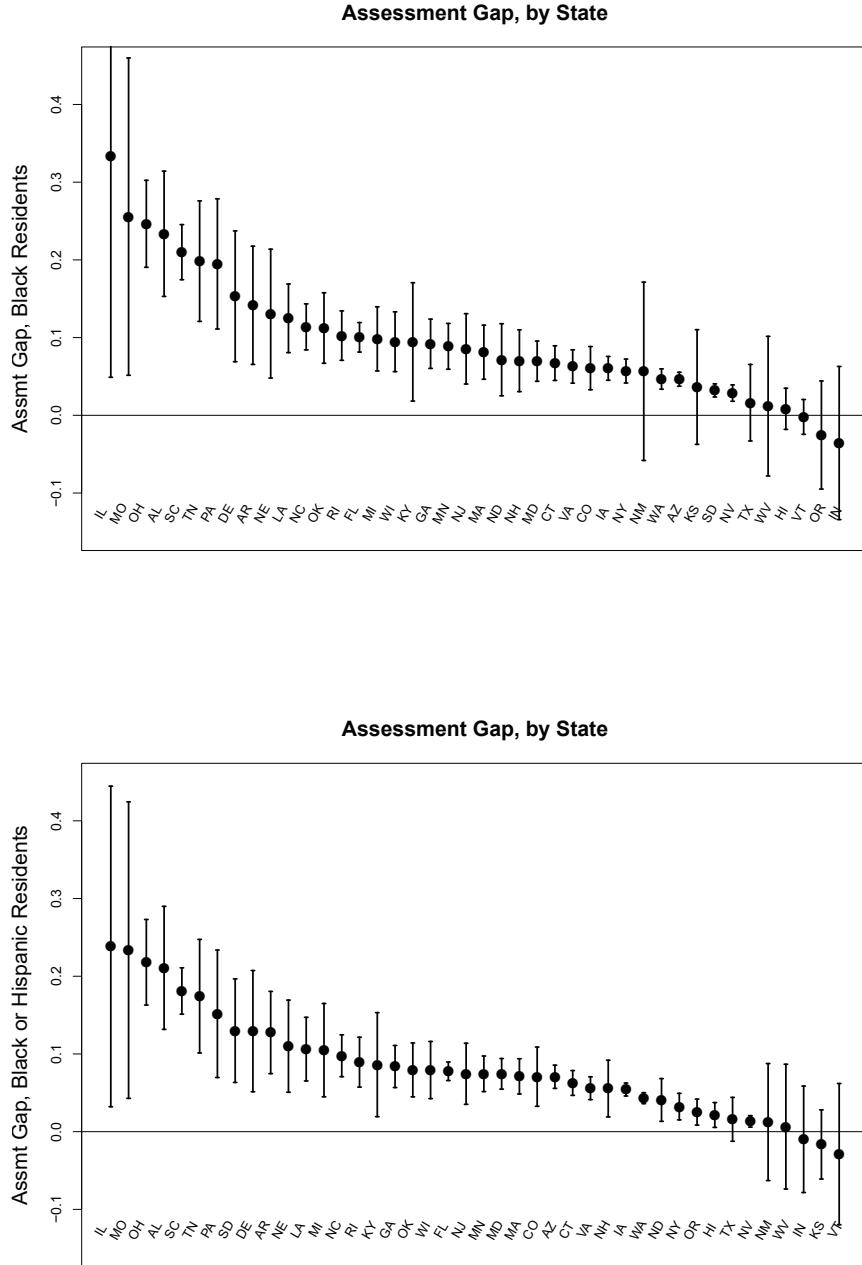
Note: This figure shows two examples to illustrate how we form taxing jurisdictions. Panel A shows a stylized example with 3 governments: a county (the large rectangle) which fully contains a city and a school district. The latter two units of government are not spatially coincident. This spatial overlay generates 4 distinct jurisdictions. Panel B presents an example with two governments: the county is again the large rectangle, and a city is entirely contained within the left (blue) portion of the county. In this example, we assume that the county is targeting a 40% assessment ratio, but realizes 50% for every home in the blue region, and realizes 20% for every home in the green region.

Figure 4: 12-Government Network in Texas



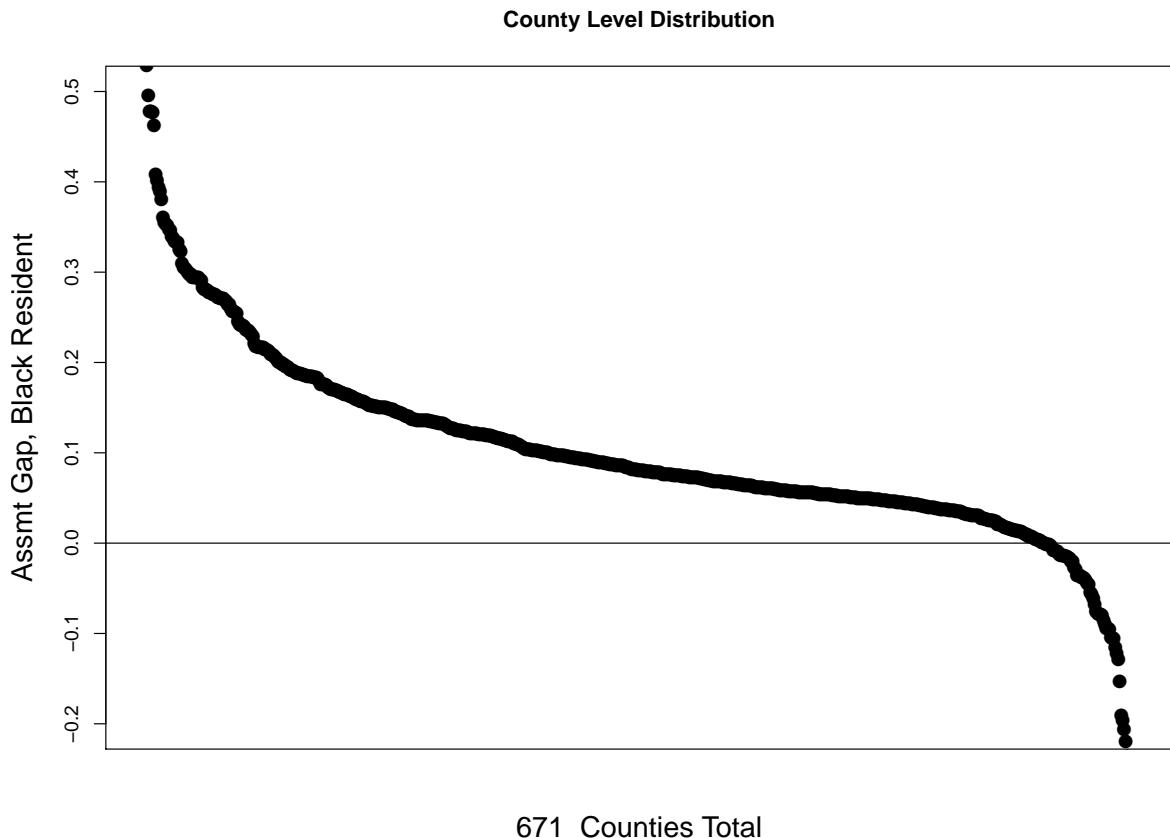
Note: This figure shows the spatial overlay of 12 different local government units in Texas. Some units are proper subsets, and thus fewer than 12 colors are evident in the figure at right. All 12 are listed at upper right. They include “standard” local governments: a county (Harris) and a city (Houston) plus two independent school districts. In addition, there are a range of entities which are related to municipal utilities or economic development initiatives. Each entity listed may, or may not, levy a property tax. Our empirical strategy generates no bias by including an entity as a taxing unit even if it does not, in fact, levy a tax in any particular year. Each unique overlapping combination of these units defines a taxing jurisdiction.

Figure 5: State Level Estimates of Assessment Gap



Note: These graphs show state-level estimates of the assessment gap. For every state with at least 500 observations, we regress log assessment ratio on a jurisdiction-year fixed effect and categorical variables for race and ethnicity. The top graph plots the estimated coefficient for black mortgage holders, along with a 95% confidence interval. The reference group is non-Hispanic white residents. Standard errors in the underlying regressions are clustered at the jurisdiction level.

Figure 6: County Level Estimates of Assessment Gap



Note: These graphs show county-level estimates of the assessment gap for black residents. For every county with at least 500 observations, we regress log assessment ratio on a jurisdiction-year fixed effect and categorical variables for race and ethnicity. We have sufficient data in 686 counties. We plot the estimated coefficient. For visual clarity, we do not include confidence intervals. Point estimates are positive and significant at 5% in 391 counties, positive and insignificant in 219 counties, negative and insignificant in 53 counties, and negative and significant at 5% in 8 counties. The reference group is non-Hispanic white residents. Standard errors in the underlying regressions are clustered at the jurisdiction level.

Figure 7: Philadelphia County/City: Demographic Heatmap

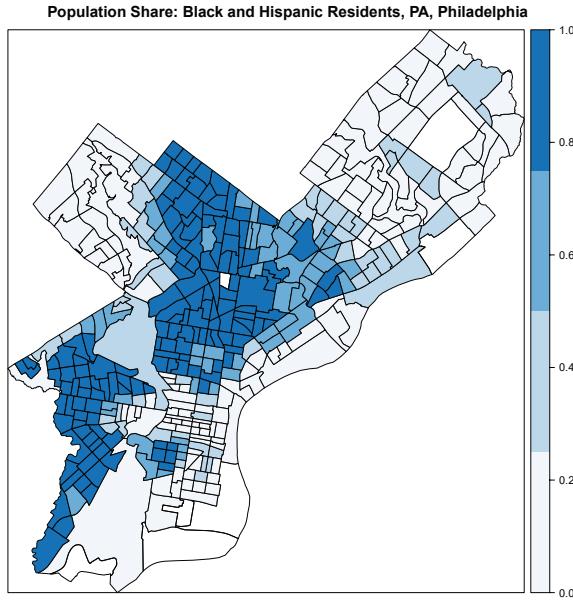
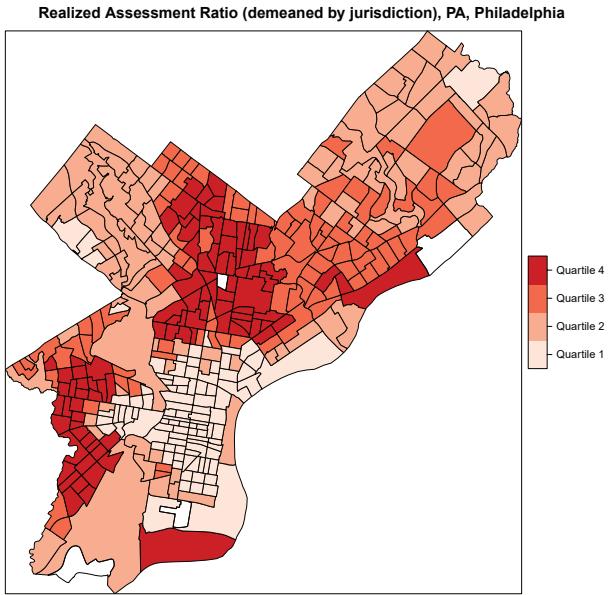


Figure 8: Philadelphia County/City: Assessment Ratio Heatmap



Note: Figure 7 plots the tract-level share of black and Hispanic residents in Philadelphia, PA using data from the American Community Survey. Tracts having a higher share of black or Hispanic residents are in darker blue. Figure 8 shows variation in realized tract-level assessment ratios computed from ATTOM. Realized log assessment ratios are residualized by jurisdiction-year, and then averaged by tract. The result, which we plot in quartiles, is an average proportional deviation from jurisdiction-mean by tract. After absorbing the jurisdiction-year means, an equitable property tax would imply no remaining variation in assessment ratio. Properties within darker red tracts have proportionately greater assessments relative to market price. Because tax bills are computed based on assessments, this mechanically represents a higher tax burden in these areas.

Figure 9: Cook County: Demographic Heatmap

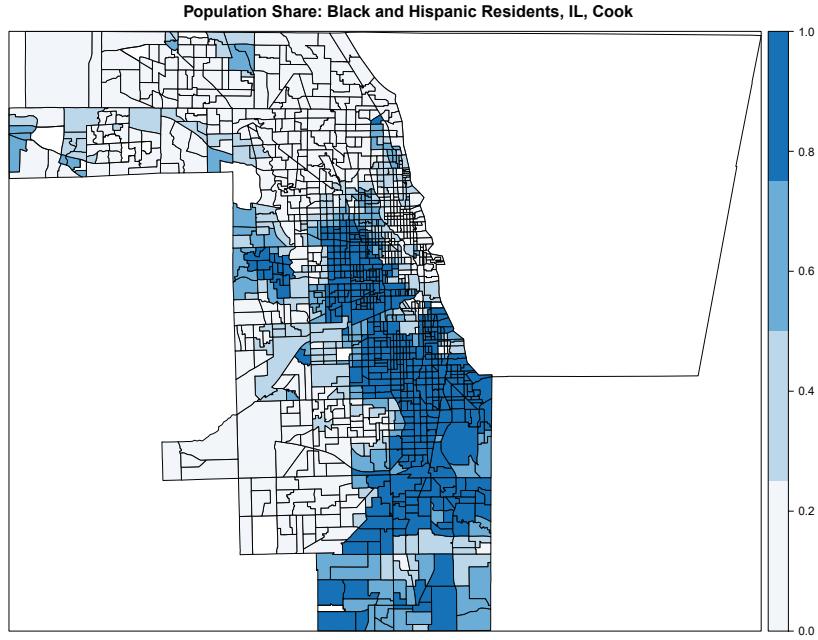
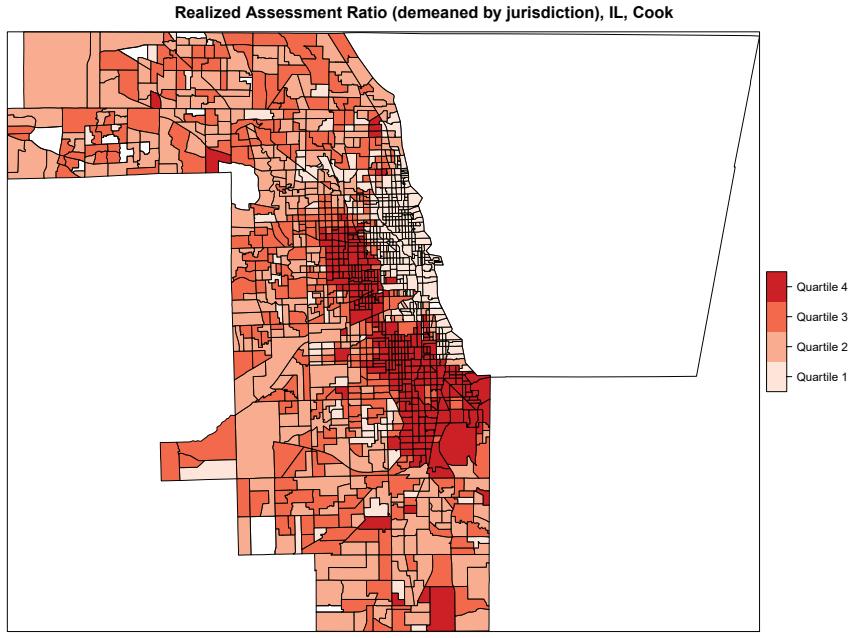


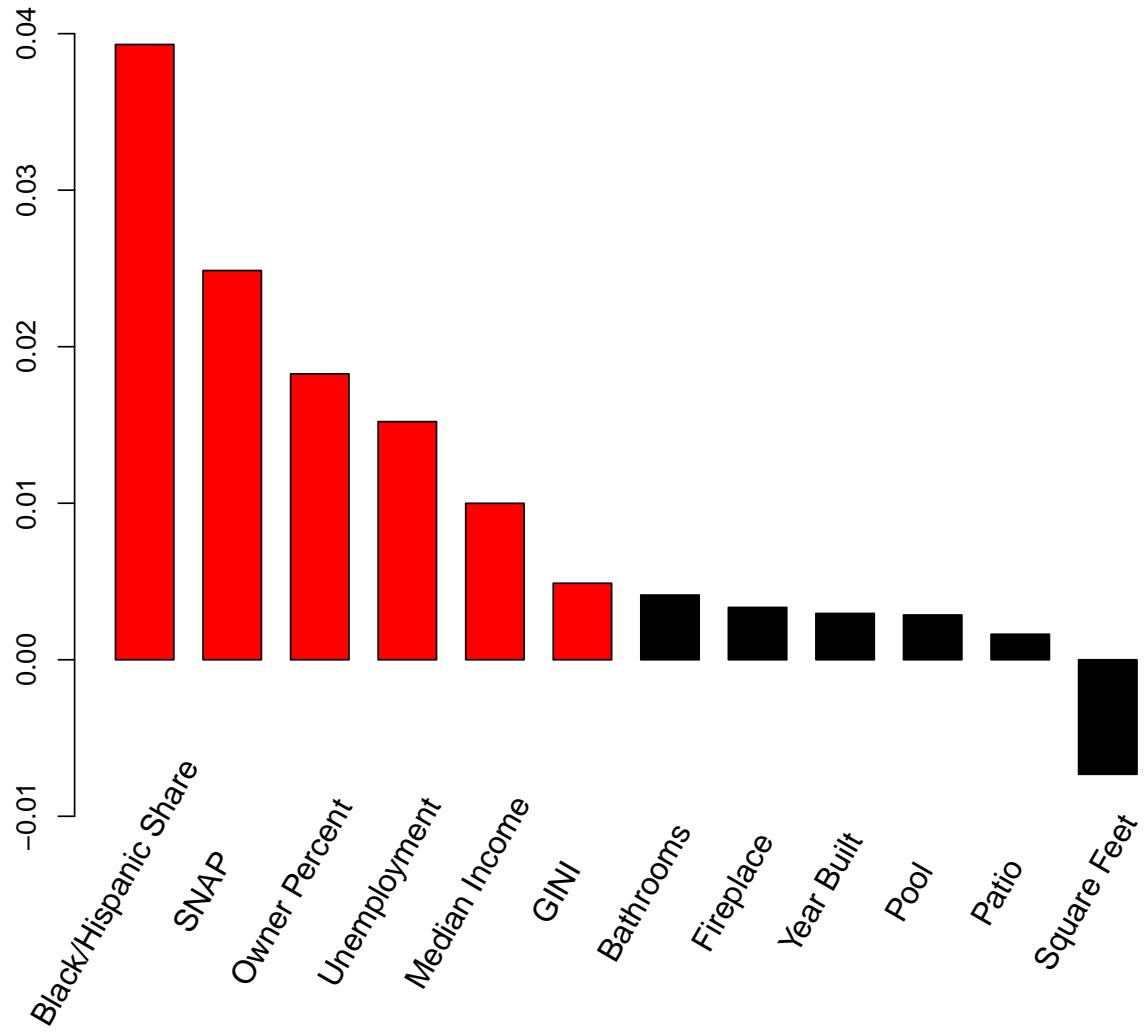
Figure 10: Cook County: Assessment Ratio Heatmap



Note: Figure 9 plots the tract-level share of black and Hispanic residents in Cook County, IL using data from the American Community Survey. Tracts having a higher share of black or Hispanic residents are in darker blue. Figure 10 shows variation in realized tract-level assessment ratios computed from ATTOM. Realized log assessment ratios are residualized by jurisdiction-year, and then averaged by tract. The result, which we plot in quartiles, is an average proportional deviation from jurisdiction-mean by tract. After absorbing the jurisdiction-year means, an equitable property tax would imply no remaining variation in assessment ratio. Properties within darker red tracts have proportionately greater assessments relative to market price. Because tax bills are computed based on assessments, this mechanically represents a higher tax burden in these areas.

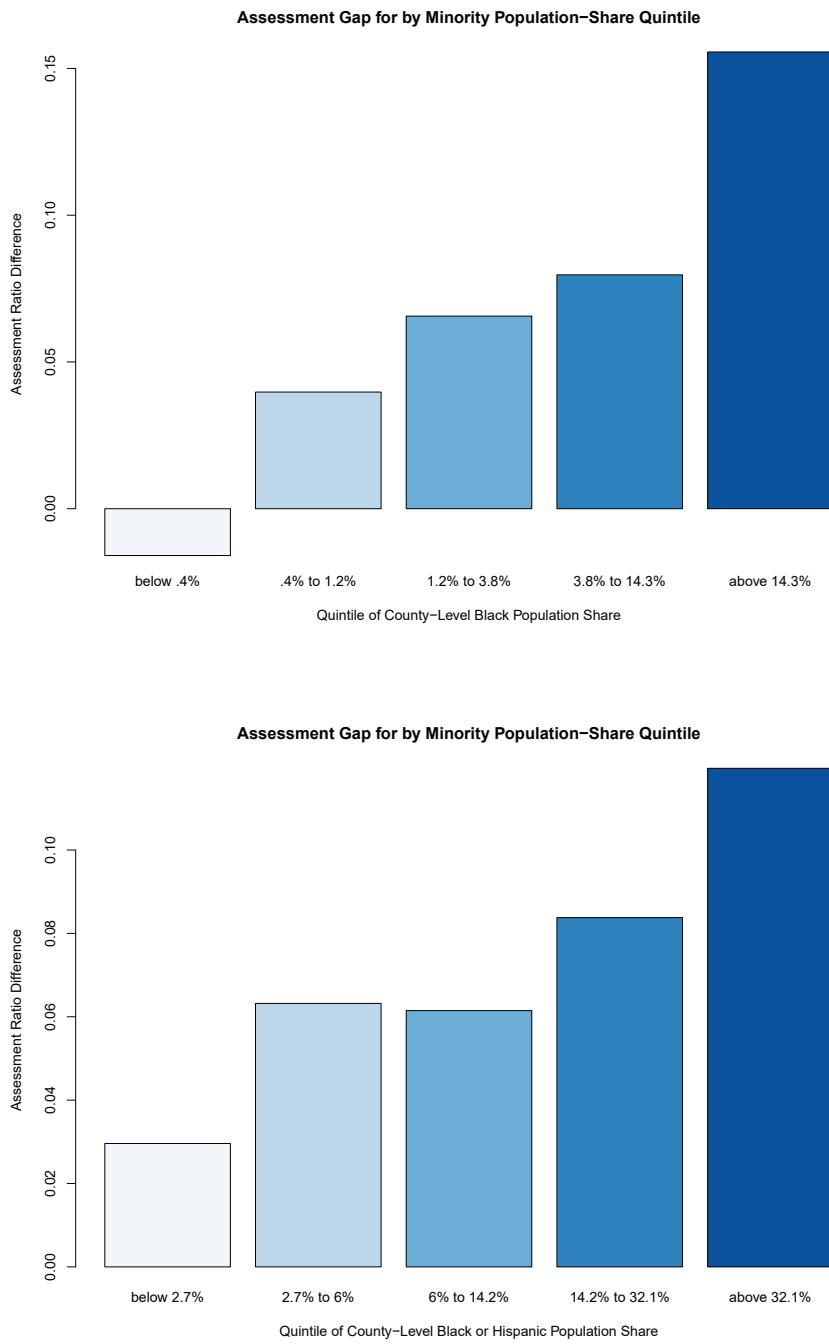
Figure 11: Hedonic Models: Mismatch

Implied Elasticity of Assessment Ratio to 1 SD Shift



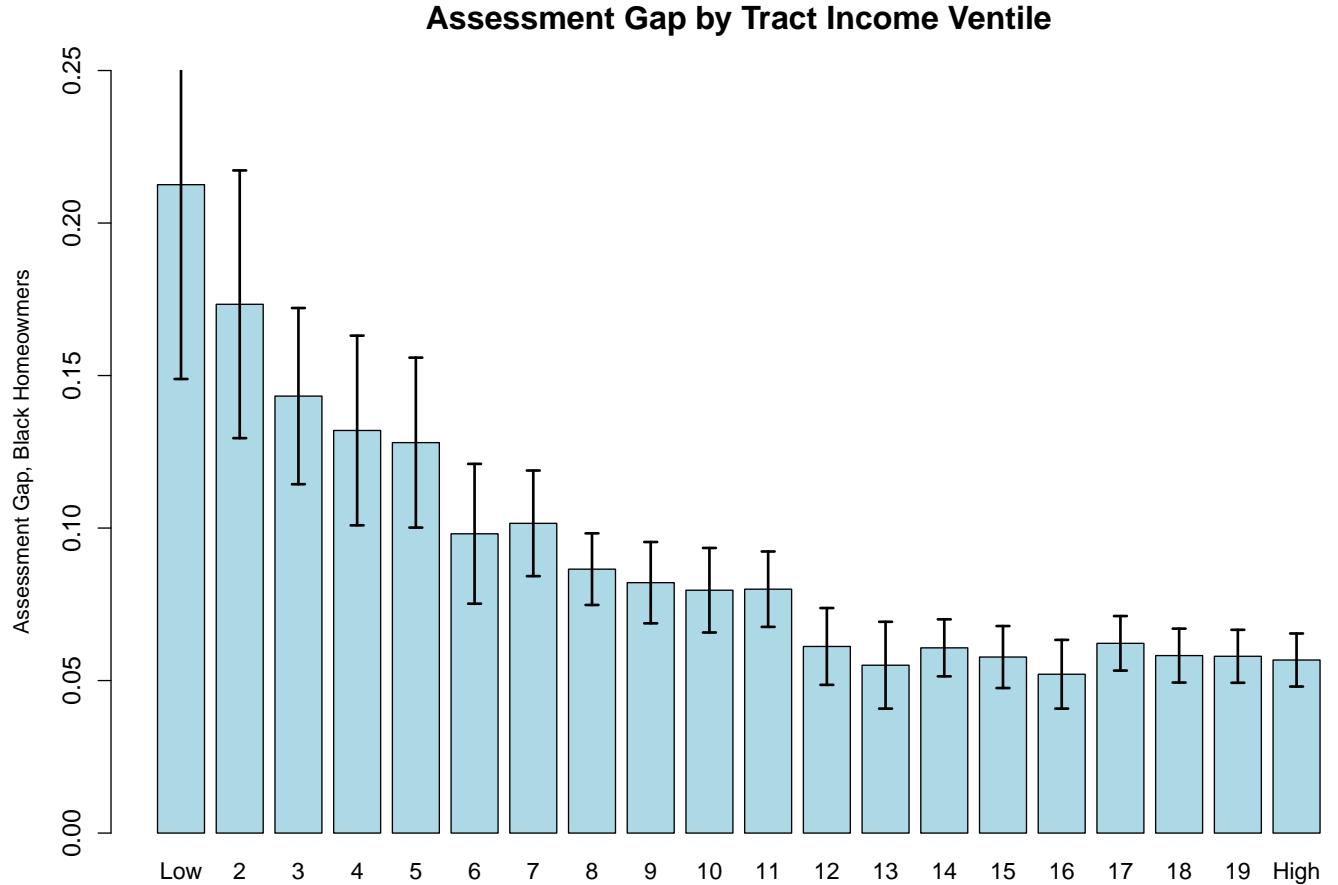
Note: Each bar in this figure plots the difference between two estimated hedonic prices: one estimated from a model with market values as the dependent variable, and one from a model with assessment values as the dependent variable. Otherwise, the two hedonic models are identical: all regressors are the same. Both market values and assessed values are logged in the underlying models, so the difference between the two estimated hedonic prices represents a proportional shift in the assessment ratio that arises from a one standard-deviation shift in the underlying variable. Bars in red are tract-level characteristics. Bars in black are property-level characteristics. A bar at zero would denote that the market-hedonic is the same as the assessment hedonic price. Larger bars signify a greater disconnect between market-hedonics and assessment-hedonics. Finally, bars above zero denote that estimated *market* hedonic prices are greater in (absolute) magnitude than assessed hedonic prices. Bars below zero denote that the assessment hedonic price is larger. Table 6 shows the estimated prices which underlie this figure.

Figure 12: Sample Split by County-Level Minority Population Share



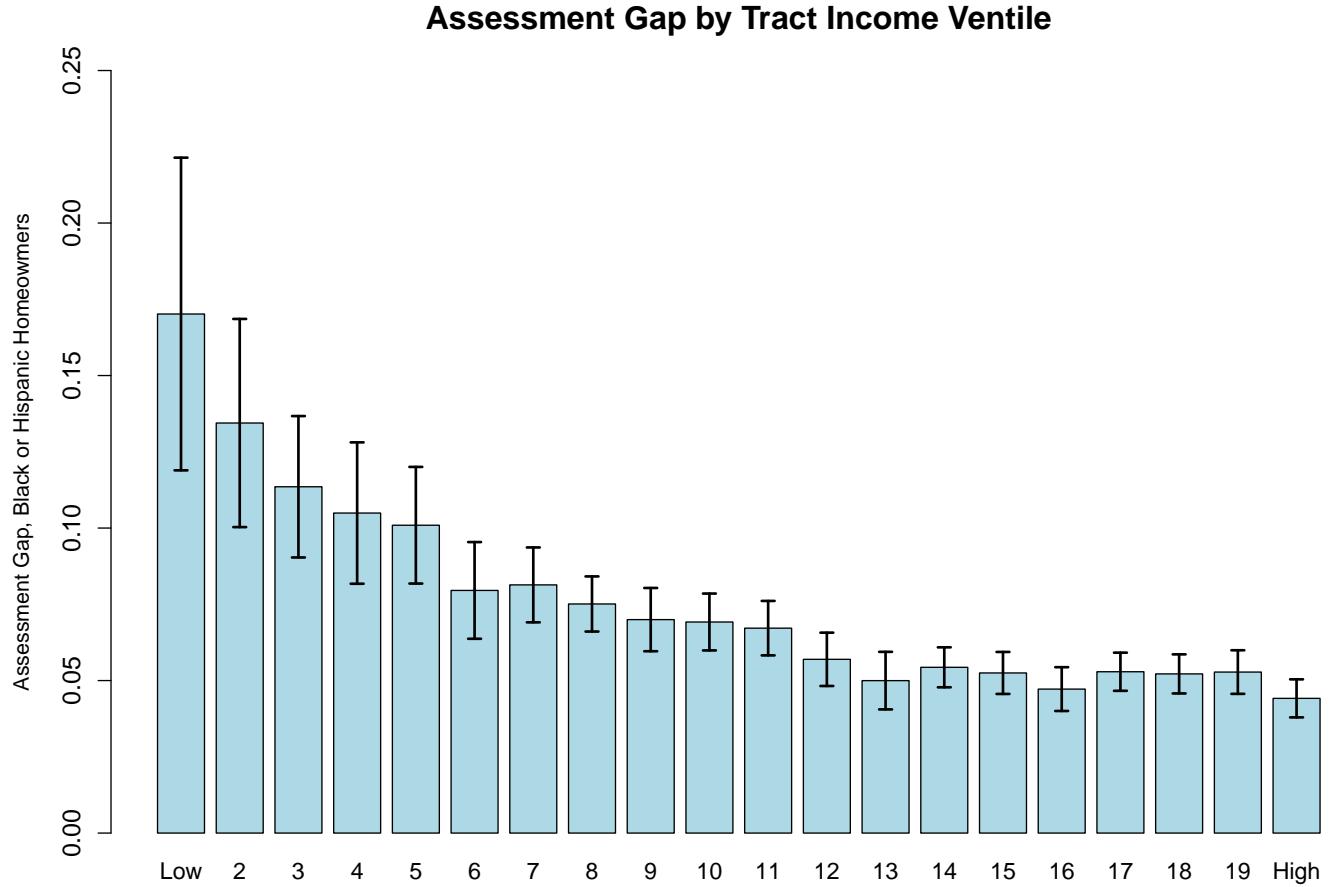
Note: These graphs show results from estimating the assessment gap in sub-samples by minority population share at the county level. We split the sample into quintiles by on average county black or black and Hispanic population share between 2005 to 2016. The quintile range is reflected below each bar. The regression output underlying this table is shown in Table 11.

Figure 13: Sample Split by Census Tract Median Income



Note: This figure shows the result of estimating the assessment gap for black homeowners in sub-samples by median census tract income. Each property's assessment ratio is residualized on a jurisdiction-year fixed effect. We then graph the mean difference between assessment ratio residuals for black homeowners and white homeowners within each of 20 quantiles based on ACS 5-year estimates of median tract-level income. "Low" denotes tracts having the lowest median income, and "high" denotes tracts with the highest median income. A 95% confidence interval is shown for each income quantile.

Figure 14: Sample Split by Census Tract Median Income



Note: This figure shows the result of estimating the assessment gap for black or Hispanic homeowners in sub-samples by median census tract income. Each property's assessment ratio is residualized on a jurisdiction-year fixed effect. We then graph the mean difference between assessment ratio residuals for black/Hispanic homeowners and white homeowners within each of 20 quantiles based on ACS 5-year estimates of median tract-level income. "Low" denotes tracts having the lowest median income, and "high" denotes tracts with the highest median income. A 95% confidence interval is shown for each income quantile.

Table 1: Baseline Assessment Gap

	log(Assessment) - log(Market)		
	(1)	(2)	(3)
Black Mortgage Holder	0.1266*** (0.0150)		
Black or Hispanic Mortgage Holder		0.0984*** (0.0106)	
Other Nonwhite Mortgage Holder			0.0278*** (0.0016)
Fixed Effects	Jurisd-Year	Jurisd-Year	Jurisd-Year
No. Clusters	37723	37723	37723
Observations	6,987,915	6,987,915	6,987,915
R ²	0.8798	0.8798	0.8798

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows our baseline findings of a racial assessment gap. We regress the log assessment ratio on a jurisdiction-year fixed effect and on categorical groupings by racial and ethnic identity. In all columns, the reference group is non-Hispanic white residents, and for clarity coefficients for groups not being considered in a given column are not reported. The estimates in this table reflect an assessment ratio differential for the given grouping of minority residents relative to non-Hispanic white residents. Standard errors are clustered at the jurisdiction level.

Table 2: Individual Race Effect: by Tract

	log(Assessment) - log(Market)		
	(1)	(2)	(3)
Black Mortgage Holder	0.0640*** (0.0020)		
Black or Hispanic Mortgage Holder		0.0530*** (0.0015)	
Other Nonwhite Mortgage Holder			0.0198*** (0.0006)
Fixed Effects	Jurisd-Tract-Yr	Jurisd-Tract-Yr	Jurisd-Tract-Yr
No. Clusters	37723	37723	37723
Observations	6,987,915	6,987,915	6,987,915
R ²	0.9005	0.9005	0.9005

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows the within-tract portion of the assessment gap. We regress the log assessment ratio on a jurisdiction-tract-year fixed effect and on categorical groupings by racial and ethnic identity. In all columns, the reference group is non-Hispanic white residents, and for clarity coefficients for groups not being considered in a given column are not reported. The estimates in this table reflect an assessment ratio differential for the given grouping of minority residents relative to non-Hispanic white residents. Standard errors are clustered at the jurisdiction level.

Table 3: Individual Race Effect: by Block Group

	log(Assessment) - log(Market)		
	(1)	(2)	(3)
Black Mortgage Holder	0.0588*** (0.0019)		
Black or Hispanic Mortgage Holder		0.0485*** (0.0014)	
Other Nonwhite Mortgage Holder			0.0190*** (0.0007)
Fixed Effects	Jurisd-BG-Yr	Jurisd-BG-Yr	Jurisd-BG-Yr
No. Clusters	37723	37723	37723
Observations	6,987,915	6,987,915	6,987,915
R ²	0.9166	0.9166	0.9166

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows the within-block group portion of the assessment gap. We regress the log assessment ratio on a jurisdiction-year-block group fixed effect and on categorical groupings by racial and ethnic identity. In all columns, the reference group is non-Hispanic white residents, and for clarity coefficients for groups not being considered in a given column are not reported. The estimates in this table reflect an assessment ratio differential for the given grouping of minority residents relative to non-Hispanic white residents. Standard errors are clustered at the jurisdiction level.

Table 4: Race and Demographic Shares

	log(Assessment) - log(Market)		
	(1)	(2)	(3)
Black Mortgage Holder	0.079*** (0.004)		
Black Share	0.299*** (0.046)		
Black or Hispanic Mortgage Holder		0.067*** (0.003)	
Black or Hispanic Share		0.277*** (0.042)	
Other Nonwhite Mortgage Holder			0.029*** (0.002)
Other Nonwhite Share			-0.139* (0.083)
Fixed Effects	Jurisd-Year	Jurisd-Year	Jurisd-Year
No. Clusters	37679	37679	37679
Observations	6,944,439	6,944,439	6,944,439
R ²	0.881	0.881	0.880

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table augments our baseline assessment gap findings in Table 1 with one measure of spatial variation: tract-level demographic shares. We regress the log assessment ratio on a jurisdiction-year fixed effect, categorical groupings by racial and ethnic identity, and tract-level demographic shares from the American Community Survey. In all columns, the reference group for mortgage holder race and ethnicity is non-Hispanic white residents, and for clarity other mortgage holder coefficients are not reported. The mortgage holder coefficients in this table reflect an assessment ratio differential for the given grouping of minority residents relative to non-Hispanic white residents. The share coefficients represent additional variation in the assessment ratio that correlates with demographic composition of the surrounding tract, holding mortgage holder race fixed. Standard errors are clustered at the jurisdiction level.

Table 5: All Neighborhood Correlates

	log(Assessment) - log(Market)		
	(1)	(2)	(3)
Black Share	0.027*** (0.005)		
Black or Hispanic Share		0.035*** (0.006)	
Other Nonwhite Share			-0.005 (0.004)
Median HH Income	-0.021*** (0.005)	-0.015*** (0.004)	-0.024*** (0.004)
Unemployment	0.015*** (0.004)	0.017*** (0.004)	0.020*** (0.005)
SNAP Assistance	0.033*** (0.004)	0.030*** (0.003)	0.040*** (0.005)
Owner Percentage	0.021*** (0.004)	0.020*** (0.004)	0.022*** (0.004)
GINI Coef	-0.011*** (0.002)	-0.009*** (0.002)	-0.012*** (0.002)
Median Age	0.003* (0.002)	0.008*** (0.003)	0.002 (0.002)
Homeowner Race Coef	0.077	0.065	0.023
Fixed Effects	Jurisd-Year	Jurisd-Year	Jurisd-Year
No. Clusters	37679	37679	37679
Observations	6,944,439	6,944,439	6,944,439
R ²	0.881	0.881	0.881

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table augments our baseline assessment gap findings in Table 1 with several measures of spatial characteristics. All regressors are tract-level variables from the American Community Survey 5-year estimates. Standard errors are clustered at the jurisdiction level. We continue to hold homeowner race fixed in this regression: those coefficients are reported in the first line of notes immediately under the estimated coefficients. Standard errors are clustered at the jurisdiction level.

Table 6: Hedonic Prices

	Market (1)	Assessment (2)	Market (3)	Assessment (4)
Black Share	-0.092*** (0.004)	-0.056*** (0.004)		
Black or Hispanic Share			-0.117*** (0.006)	-0.078*** (0.005)
Median HH Income	0.157*** (0.008)	0.144*** (0.008)	0.145*** (0.008)	0.135*** (0.008)
Unemployment	-0.027*** (0.003)	-0.013*** (0.002)	-0.030*** (0.004)	-0.015*** (0.002)
SNAP Share	-0.089*** (0.006)	-0.061*** (0.004)	-0.075*** (0.006)	-0.050*** (0.004)
Owner Share	-0.049*** (0.005)	-0.032*** (0.003)	-0.053*** (0.005)	-0.035*** (0.004)
GINI	0.066*** (0.004)	0.059*** (0.004)	0.058*** (0.004)	0.053*** (0.004)
Square Feet	0.256*** (0.029)	0.264*** (0.030)	0.256*** (0.029)	0.264*** (0.030)
Bathrooms	0.107*** (0.017)	0.103*** (0.017)	0.107*** (0.017)	0.103*** (0.017)
Year Built	0.031*** (0.003)	0.028*** (0.003)	0.030*** (0.003)	0.028*** (0.003)
Other Attributes	Y	Y	Y	Y
Fixed Effects	Jurisd-Year	Jurisd-Year	Jurisd-Year	Jurisd-Year
No. Clusters	26152	26152	26152	26152
Observations	4,877,658	4,877,658	4,877,658	4,877,658
R ²	0.773	0.942	0.773	0.942

Note:

*p<0.1; **p<0.05; ***p<0.01

Not shown: coefficients on indicators for patio, pool, and fireplace

Note: This table reports estimated hedonic prices from two separate hedonic models. The first model uses (log) market as the dependent variable. These estimates are reported in columns 1 and 3. The second model uses (log) assessed values as the dependent variable. These estimates are reported in columns 2 and 4. Otherwise, the two hedonic models are identical: all regressors are the same. The table omits estimated coefficients for indicator variables stating whether a property has a patio, pool, or fireplace. Standard errors are clustered at the jurisdiction level. Figure 11 shows the difference between attribute-coefficients graphically.

Table 7: Cook County Appeals

	Dependent Variable:		
	Appeal	Win Appeal	Reduction
	(1)	(2)	(3)
Black Mortgage Holder	-0.840*** (0.083)	-2.193*** (0.354)	-0.480*** (0.117)
Baseline Rate	14.6	67.4	12.0
Fixed Effects	BG-Year	BG-Year	BG-Year
No. Clusters	3954	3933	3893
Observations	4,076,655	694,553	476,368
R ²	0.383	0.415	0.442

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table uses administrative microdata on property tax appeals in Cook County. The first column shows unconditional propensity to appeal. Column 2 conditions on a homeowner having filed an assessment appeal. Column 3 conditions on a successful appeal. In columns 1 and 2, the dependent variable is a binary indicator. In column 3, the dependent variable is the reduction amount divided by the proposed assessment. Fixed effects across all columns are at the block-group-year level. Standard errors are clustered at the block group level. The baseline rates for (i) appeal propensity, (ii) winning appeal, and (iii) reduction conditional on a successful appeal are reported in the first line below the estimates. Coefficients and baseline rates are reported as percents.

Table 8: Cook County Appeals

	Dependent Variable:		
	Appeal	Win Appeal	Reduction
	(1)	(2)	(3)
Black or Hispanic Mortgage Holder	-0.982*** (0.068)	-1.993*** (0.245)	-0.258*** (0.074)
Baseline Rate	14.6	67.4	12.0
Fixed Effects	BG-Year	BG-Year	BG-Year
No. Clusters	3954	3933	3893
Observations	4,076,655	694,553	476,368
R ²	0.383	0.415	0.443

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table uses administrative microdata on property tax appeals in Cook County. The first column shows unconditional propensity to appeal. Column 2 conditions on a homeowner having filed an assessment appeal. Column 3 conditions on a successful appeal. In columns 1 and 2, the dependent variable is a binary indicator. In column 3, the dependent variable is the reduction amount divided by the proposed assessment. Fixed effects across all columns are at the block-group-year level. Standard errors are clustered at the block group level. The baseline rates for (i) appeal propensity, (ii) winning appeal, and (iii) reduction conditional on a successful appeal are reported in the first line below the estimates. Coefficients and baseline rates are reported as percents.

Table 9: Effect of Racial Ownership on Assessments

	Assessments			
	Growth		Levels	
	(1)	(2)	(3)	(4)
Black Mortgage Holder	0.0711* (0.0386)		0.2917*** (0.0415)	
Black or Hispanic Mortgage Holder		0.4103*** (0.0255)		0.7923*** (0.0274)
Fixed Effects	Two-Way	Two-Way	Two-Way	Two-Way
No. Clusters	12268641	12268641	12268641	12268641
Observations	54,970,191	54,970,191	54,970,191	54,970,191
R ²	0.6925	0.6925	0.9910	0.9910

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows the results of a generalized difference-in-differences estimation. The dependent variable is logged assessment value. Every home in this sample is transacted at least once. Fixed effects are two-way: property and tract-year. In columns 1 and 2, the dependent variable is growth rates (log difference in assessed value). In columns 3 and 4, the dependent variable is the logged assessment. Standard errors are clustered at the property level.

Table 10: Sample Split by Racial Attitudes

	Assessment Value / Market Value				
	Baseline		By Media Market		By State
	(1)	(2)	(3)	(4)	(5)
Black Mortgage Holder	0.128*** (0.015)				
Black, High Animus		0.150*** (0.022)	0.070*** (0.003)	0.145*** (0.011)	0.076*** (0.003)
Black, Low Animus		0.084*** (0.008)	0.055*** (0.002)	0.106*** (0.033)	0.049*** (0.002)
Fixed Effects	Jurisd-Yr	Jursid-Yr	Jurisd-Tract-Yr	Jurisd-Yr	Jursid-Tract-Yr
No. Clusters	37106	37106	37106	37106	37106
Observations	6,856,585	6,856,585	6,856,585	6,856,585	6,856,585
R ²	0.881	0.881	0.902	0.881	0.902

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows results of using the measures of racial animus described in [Stephens-Davidowitz \(2014\)](#) to split our sample into regions of above- and below-median prejudice. Column 1 shows baseline results before splitting the sample. Columns 2 and 3 use a media-market measure of animus. We use a Nielsen crosswalk to associate media markets with individual counties. Columns 4 and 5 use a state-level measure of animus. For each measure, the first result (column 2 or 4) shows the overall assessment gap. The second result shows the homeowner effect estimated within jurisdiction-tract-year. For all specifications, standard errors are clustered at the jurisdiction level.

Table 11: Sample Split by County-Level Minority Population Share

Panel A

	Assessment Value / Market Value				
	Quintile of County-Level Minority Population Share				
	(1)	(2)	(3)	(4)	(5)
Black Mortgage Holder	-0.016 (0.054)	0.040*** (0.007)	0.066*** (0.004)	0.080*** (0.006)	0.156*** (0.022)
Fixed Effects	Jurisd-Yr	Jurisd-Yr	Jurisd-Yr	Jurisd-Yr	Jurisd-Yr
No. Clusters	2008	6491	9490	12813	6323
Observations	53,919	405,323	909,640	3,114,742	2,372,961
R ²	0.856	0.938	0.906	0.888	0.850

Note:

*p<0.1; **p<0.05; ***p<0.01

Panel B

	Assessment Value / Market Value				
	Quintile of County-Level Minority Population Share				
	(1)	(2)	(3)	(4)	(5)
Black or Hispanic Mortgage Holder	0.030** (0.014)	0.063*** (0.006)	0.061*** (0.003)	0.084*** (0.006)	0.120*** (0.019)
Fixed Effects	Jurisd-Yr	Jurisd-Yr	Jurisd-Yr	Jurisd-Yr	Jurisd-Yr
No. Clusters	3215	5989	10998	12089	4843
Observations	73,243	295,057	1,433,767	2,796,141	2,258,377
R ²	0.819	0.786	0.858	0.879	0.882

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: Each panel shows the results from estimating the assessment gap on sub-samples based on county-level demographics. For Panel A, we split our baseline sample into quintiles by average county black population share. In Panel B the sample is split by black or Hispanic population share. In each panel, column 1 shows the estimated assessment gap within the lowest minority-population quintile, and column 5 shows results for the highest quintile. Regressions are run separately rather than pooled. We include jurisdiction-year fixed effects in all specifications. Standard errors are clustered at the jurisdiction level.

Table 12: Assessment Ratio Pass Through to Tax Bill

	Effective Tax Rate - Year of Sale (%)		
	Tax Bill	Tax Bill	Before Exemptions
	(1)	(2)	(3)
All Mortgage Holders	0.9913*** (0.0039)		
White Mortgage Holder		0.9925*** (0.0037)	0.8569*** (0.0128)
Black or Hispanic Mortgage Holder		0.9857*** (0.0056)	0.8517*** (0.0131)
Other Nonwhite Mortgage Holder		0.9892*** (0.0040)	0.8536*** (0.0131)
Fixed Effects	Jurisd-Year	Jurisd-Year	Jurisd-Year
No. Clusters	26371	26371	26371
Observations	3,373,164	3,373,164	3,373,164
R ²	0.9191	0.9192	0.7672

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows the results of regressing log effective tax rate on log assessment ratio. Column 1 presents estimates for all homeowners. Columns 2 and 3 show a breakdown by racial and ethnic grouping. Results for black homeowners alone are very similar to those reported here. In columns 1 and 2, the dependent variable is an effective rate formed using the actual tax bill reported in the ATTOM dataset. Column 3 computes a pre-exemption effective rate by adding reported exemptions back to the reported tax bill. The effective rate is computed by using the tax bill reported in the same year as the sale. All specifications use jurisdiction-year fixed effects. Standard errors are clustered at the jurisdiction level.

Table 13: Effective Tax Rate, Sale Year

	Effective Tax Rate - In Sale Year (%)					
	Tax Bill	Before Exemptions	Tax Bill	Before Exemptions	Tax Bill	Before Exemptions
	(1)	(2)	(3)	(4)	(5)	(6)
Black Mortgage Holder	14.8834*** (1.9459)	12.2187*** (2.0551)				
Black or Hispanic Mortgage Holder			11.3977*** (1.4335)	8.0480*** (1.5783)		
Other Nonwhite Mortgage Holder					3.2118*** (0.2287)	2.0736*** (0.2737)
Jurisd-Year FE	Y	Y	Y	Y	Y	Y
Other Controls	N	N	N	N	N	N
No. Clusters	26371	26371	26371	26371	26371	26371
Observations	3,373,164	3,373,164	3,373,164	3,373,164	3,373,164	3,373,164
R ²	0.6803	0.6481	0.6802	0.6478	0.6802	0.6478

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table repeats our baseline analysis in Table 1, but uses effective tax rate as the dependent variable instead of assessment ratio. Coefficients are percentages. For each racial and ethnic grouping, we present two sets of results. In odd columns, we show results using an effective rate computed using the observed tax bill and observed market value in the same year. Because the observed tax bill is potentially net of a wide range of exemptions, we also compute a before-exemption effective tax rate, by adding observed exemptions to the observed tax bill, and then dividing by market value. We trim any observation above a calculated effective tax rate of 25% both before and net of exemptions. We believe this to be a conservative choice as 25% is far higher than any property tax rate of which we are aware (the national median is approximately 1.4%), and is more likely than not to be a data error. All specifications use jurisdiction-year fixed effects to hold constant the level of intended taxation. Standard errors are clustered at the jurisdiction level.

Table 14: Racial Differential in Transacted Prices

	Proportional Realized Price Difference		
	(1)	(2)	(3)
Black Seller	0.022*** (0.002)		
Black or Hispanic Seller		0.033*** (0.002)	
Other Non-White Seller			0.005*** (0.001)
Fixed Effects	Jurisd-B.G.-Yr	Jurisd-B.G.-Yr	Jurisd-B.G.-Yr
No. Clusters	18984	18984	18984
Observations	2,196,003	2,196,003	2,196,003
R ²	0.801	0.802	0.802

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows results from regressing the log difference of realized market price and predicted market price on a block-group-year fixed effect and categorical groupings by racial and ethnic identity. In all columns, the reference group is non-Hispanic white residents, and for clarity coefficients for groups not being considered in a given column are not reported. The estimates in this table reflect a racial differential in transaction prices net of predicted price. The predicted price is generated using zip-code level home price indexes. Standard errors are clustered at the jurisdiction level.

Table 15: Synthetic Assessments Using Zip Code HPIs

	log(Assessment) - log(Market)					
	Real Assessments			Synthetic Assessments		
	(1)	(2)	(3)	(4)	(5)	(6)
Black Mortgage Holder	0.144*** (0.015)			-0.041*** (0.003)		
Black or Hispanic Mortgage Holder		0.110*** (0.011)			-0.051*** (0.003)	
Other Nonwhite Mortgage Holder			0.031*** (0.002)			-0.007*** (0.001)
Jurisd-Year FE	Y	Y	Y	Y	Y	Y
Other Controls	N	N	N	N	N	N
No. Clusters	18853	18853	18853	18853	18853	18853
Observations	2,135,943	2,135,943	2,135,943	2,135,943	2,135,943	2,135,943
R ²	0.910	0.910	0.910	0.712	0.713	0.713

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows the results from our proposed approach for correcting the assessment gap. Using the algorithm described in Section 6, we construct synthetic assessments using zip-code-level HPIs. We use Zillow's publicly available ZHVI series by zip-code. Our approach uses an initial transaction to pin down the base assessment value. At every subsequent transaction, we observe a realized assessment ratio along with our synthetically constructed assessment ratio. Columns 1–3 show that the overall assessment gap looks similar in the subset of homes to which we can apply this approach (smaller chiefly because the first transaction is not included in the analysis). Columns 4–6 show the assessment gap using our synthetic assessment ratios. All specifications include jurisdiction-year fixed effects. Standard errors are clustered at the jurisdiction level.

Table A1: Assessment Ratio Differentials in California

	Assessment Value / Market Value		
	(1)	(2)	(3)
Black Mortgage Holder	0.0413*** (0.0101)		
Black or Hispanic Mortgage Holder		0.1060*** (0.0044)	
Other Nonwhite Mortgage Holder			0.0653*** (0.0030)
Fixed Effects	Jurisd-Year	Jurisd-Year	Jurisd-Year
No. Clusters	5603	5603	5603
Observations	1,186,388	1,186,388	1,186,388
R ²	0.3816	0.3820	0.3820

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows the results of our baseline assessment gap analysis for California alone. We regress the log assessment ratio on a jurisdiction-year fixed effect and on categorical groupings by racial and ethnic identity. In all columns, the reference group is non-Hispanic white residents, and for clarity coefficients for groups not being considered in a given column are not reported. The estimates in this table reflect an assessment ratio differential for the given grouping of minority residents relative to non-Hispanic white residents. Standard errors are clustered at the jurisdiction level.

Table A2: Assessment Gap, Using Counties instead of Taxing Jurisdictions

	log(Assessment) - log(Market)		
	(1)	(2)	(3)
Black Mortgage Holder	0.1687*** (0.0187)		
Black or Hispanic Mortgage Holder		0.1356*** (0.0138)	
Other Nonwhite Mortgage Holder			0.0321*** (0.0024)
Fixed Effects	County-Year	County-Year	County-Year
No. Clusters	1982	1982	1982
Observations	6,987,915	6,987,915	6,987,915
R ²	0.8507	0.8508	0.8508

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table repeats our baseline assessment gap analysis, but uses county-year fixed effects rather than jurisdiction-year. We regress the log assessment ratio on a county-year fixed effect and on categorical groupings by racial and ethnic identity. In all columns, the reference group is non-Hispanic white residents, and for clarity coefficients for groups not being considered in a given column are not reported. The estimates in this table reflect an assessment ratio differential for the given grouping of minority residents, relative to non-Hispanic white residents. Standard errors are clustered at the county level. This specification shows that our results are not driven by the way we form jurisdictions. Our preferred specifications all use the more rigorous within-jurisdiction analysis.

Table A3: Robustness: Jurisdiction-Month-Year Fixed Effects

	log(Assessment) - log(Market)		
	(1)	(2)	(3)
Black Mortgage Holder	0.1283*** (0.0174)		
Black or Hispanic Mortgage Holder		0.0988*** (0.0124)	
Other Nonwhite Mortgage Holder			0.0282*** (0.0019)
Fixed Effects		Jurisdiction-Month-Year	
No. Clusters	37723	37723	37723
Observations	6,987,915	6,987,915	6,987,915
R ²	0.9000	0.8999	0.8999

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table repeats our baseline assessment gap analysis, but uses jurisdiction-month-year fixed effects instead of jurisdiction-year fixed effects. We regress the log assessment ratio on a jurisdiction-month-year fixed effect and on categorical groupings by racial and ethnic identity. In all columns, the reference group is non-Hispanic white residents, and for clarity coefficients for groups not being considered in a given column are not reported. The estimates in this table reflect an assessment ratio differential for the given grouping of minority residents relative to non-Hispanic white residents. Standard errors are clustered at the jurisdiction level. This specification shows that measurement error introduced by forming fixed effects with calendar years rather than (unobserved) fiscal years does not lead to meaningfully different estimates.

Table A4: Assessment Gap Using Simple Ratios

	Assessment Value / Market Value		
	(1)	(2)	(3)
Black Mortgage Holder	0.0897*** (0.0057)		
Black or Hispanic Mortgage Holder		0.0696*** (0.0039)	
Other Nonwhite Mortgage Holder			0.0208*** (0.0010)
Fixed Effects	Jurisd-Year	Jurisd-Year	Jurisd-Year
No. Clusters	37723	37723	37723
Observations	6,987,915	6,987,915	6,987,915
R ²	0.6987	0.6986	0.6986

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows the results of our baseline assessment gap analysis using simple ratios (assessment divided by market) as the dependent variable instead of logged assessment ratios. We regress the simple assessment ratio on a jurisdiction-year fixed effect and on categorical groupings by racial and ethnic identity. In all columns, the reference group is non-Hispanic white residents, and for clarity coefficients for groups not being considered in a given column are not reported. The estimates in this table are intended to show that our baseline findings are not being mechanically generated by a Jensen's inequality term arising from taking the log of the assessment ratio. We use logged assessment ratios in our preferred specifications because the target assessment ratio varies widely across jurisdictions, and we wish to estimate proportional variation rather than variation in levels. Other than showing that inequalities do not disappear when using simple ratios, the estimates in this table have little intuition. The estimates in this table are a weighted average of absolute variation around jurisdiction means ranging from 7% to 100%. It is therefore natural that the results are lower than our baseline findings in Table 1. Standard errors are clustered at the jurisdiction level.

Table A5: Effective Tax Rate, One Year Before Sale

	Effective Tax Rate - One Year Before Sale (%)					
	Tax Bill	Before Exemptions	Tax Bill	Before Exemptions	Tax Bill	Before Exemptions
	(1)	(2)	(3)	(4)	(5)	(6)
Black Mortgage Holder	15.2528*** (2.0458)	12.2586*** (2.1646)				
Black or Hispanic Mortgage Holder			11.6826*** (1.4850)	7.8133*** (1.6357)		
Other Nonwhite Mortgage Holder					3.1404*** (0.2550)	2.0352*** (0.2959)
Jurisd-Year FE	Y	Y	Y	Y	Y	Y
Other Controls	N	N	N	N	N	N
No. Clusters	26371	26371	26371	26371	26371	26371
Observations	3,373,164	3,373,164	3,373,164	3,373,164	3,373,164	3,373,164
R ²	0.6659	0.6315	0.6657	0.6312	0.6657	0.6312

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table repeats our analysis in Table 13, but uses the tax bill from the year before sale. The denominator for computing the effective tax rate remains the observed market value. Coefficients are percentages. For each racial and ethnic grouping we present two sets of results. In odd columns, we show results using an effective rate computed using the observed tax bill and observed market value. Because the observed tax bill is potentially net of a wide range of exemptions, we also compute a before-exemption effective tax rate, by adding observed exemptions to the observed tax bill and then dividing by market value. We trim any observation above a calculated effective tax rate of 25% both before and net of exemptions. We believe this to be a conservative choice as 25% is far higher than any property tax rate of which we are aware (the national median is approximately 1.4%), and is more likely than not to be a data error. All specifications use jurisdiction-year fixed effects to hold constant the level of intended taxation. Standard errors are clustered at the jurisdiction level.

Table A6: Effective Tax Rate, One Year After Sale

	Effective Tax Rate - One Year After Sale (%)					
	Tax Bill	Before Exemptions	Tax Bill	Before Exemptions	Tax Bill	Before Exemptions
	(1)	(2)	(3)	(4)	(5)	(6)
Black Mortgage Holder	13.1055*** (1.8480)	10.2602*** (1.9628)				
Black or Hispanic Mortgage Holder			9.7809*** (1.3657)	7.0178*** (1.4751)		
Other Nonwhite Mortgage Holder					2.9336*** (0.2023)	2.0251*** (0.2154)
Jurisd-Year FE	Y	Y	Y	Y	Y	Y
Other Controls	N	N	N	N	N	N
No. Clusters	26371	26371	26371	26371	26371	26371
Observations	3,373,164	3,373,164	3,373,164	3,373,164	3,373,164	3,373,164
R ²	0.7042	0.6703	0.7039	0.6701	0.7039	0.6701

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table repeats our analysis in Table 13, but uses the tax bill from the year after the sale. The denominator for computing the effective tax rate remains the observed market value. Coefficients are percentages. For each racial and ethnic grouping we present two sets of results. In odd columns, we show results using an effective rate computed using the observed tax bill and observed market value. Because the observed tax bill is potentially net of a wide range of exemptions, we also compute a before-exemption effective tax rate, by adding observed exemptions to the observed tax bill, and then dividing by market value. We trim any observation above a calculated effective tax rate of 25% both before and net of exemptions. We believe this to be a conservative choice as 25% is far higher than any property tax rate of which we are aware (the national median is approximately 1.4%), and is more likely than not to be a data error. All specifications use jurisdiction-year fixed effects to hold constant the level of intended taxation. Standard errors are clustered at the jurisdiction level.

Table A7: Synthetic Assessments, Stopping Growth in January Each Year

	log(Assessment) - log(Market)					
	Real Assessments			Synthetic Assessments		
	(1)	(2)	(3)	(4)	(5)	(6)
Black Mortgage Holder	0.144*** (0.015)			-0.040*** (0.003)		
Black or Hispanic Mortgage Holder		0.110*** (0.011)			-0.049*** (0.003)	
Other Nonwhite Mortgage Holder			0.031*** (0.002)			-0.007*** (0.001)
Jurisd-Year FE	Y	Y	Y	Y	Y	Y
Other Controls	N	N	N	N	N	N
No. Clusters	18853	18853	18853	18853	18853	18853
Observations	2,135,943	2,135,943	2,135,943	2,135,943	2,135,943	2,135,943
R ²	0.910	0.910	0.910	0.692	0.693	0.693

Note:

*p<0.1; **p<0.05; ***p<0.01

Note: This table shows an alternative implementation of our proposed approach for correcting the assessment gap. The analysis in Table 15 uses constructed assessments which increase with the zip-code HPI until the month of sale. In this table, we use constructed assessments which change only in January of each year. This more closely parallels the actual assessment practice of generating a single value each year. In this approach, when a sale occurs, the assessment is out of date by up to 12 months. Columns 1–3 are identical to Table 15 and show that the overall assessment gap looks similar in the subset of homes to which we can apply this approach. Columns 4–6 show the assessment gap using January-revised synthetic assessments. All specifications include jurisdiction-year fixed effects. Standard errors are clustered at the jurisdiction level.