

Algorithmic Governance and Tax Equity: Evidence from Property Assessment Reform in Cook County

Meng Liu*

April 8, 2026

Abstract

We study whether algorithmic reform can reduce inequality in residential property taxation. In Cook County's 2021 algorithmic reform, a legacy valuation system was replaced by a more flexible automated valuation model using richer property and geospatial data. Using the county's staggered triennial reassessment cycle across three regions, we estimate difference-in-differences models of the log assessed-to-sale-price ratio, focusing on price and race gradients. The reform sharply reduces long-standing price regressivity favoring the rich: relative to 2020, the price gradient flattens by about 40% in the first post-reform quarter and remains nearly unchanged after appeals, implying little erosion of equity gains. Race effects are mixed: the within-price race slope rises from a negative baseline (i.e., pro-minority), but total race regressivity becomes somewhat more negative because the reform's dominant distributional effect compresses the price gradient. Mechanism evidence suggests that legacy systems overpool unlike properties: linear models can overassess lower-value homes by pooling them with higher-value homes, while coarse location controls can underassess homes in minority neighborhoods, conditional on price, by linking them to lower-valued neighborhood averages. More flexible machine-learning models with finer geography relax both distortions, reducing price regressivity and shifting the within-price racial gradient.

*Olin Business School, Washington University in St. Louis. Email: mengl@wustl.edu

JEL Codes: H71, H20, C55, D73, R31

Keywords: algorithmic governance, machine learning, property tax, tax equity, regressivity, government AI adoption

1 Introduction

Can algorithmic reform reduce inequality in a core function of government? We study this question in the context of residential property assessment, the administrative process that determines the tax base for the property tax. The property tax is the fiscal backbone of local government, and its fairness depends on a simple *ad valorem* principle: homes of equal market value within the same taxing jurisdiction should bear equal tax burdens. Yet a growing literature shows that this principle is often violated. Lower-priced homes are systematically overassessed relative to higher-priced homes, making effective property taxation regressive with respect to wealth (Berry, 2021; Amornsiripanitch, 2020). Black and Hispanic homeowners also face higher assessment ratios and higher property tax burdens than similarly situated white homeowners, both because assessments fail to fully reflect neighborhood amenities and because downstream institutions such as appeals and exemptions are unequally navigated (Avenancio-León and Howard, 2022). These distortions matter not only for distribution, but also for the legitimacy of the fiscal state (Atuahene and Berry, 2018).

A natural response is algorithmic modernization. Many assessor offices still rely on legacy mass-appraisal systems built around coarse geographic strata, limited covariates, rigid functional forms, and opaque manual adjustments. If longstanding inequities arise because those systems fit heterogeneous housing markets poorly, then richer data and more flexible predictive technology should improve both accuracy and equity. But that conclusion is not automatic. A broad literature on algorithms in government emphasizes both promise and peril: data-driven systems may improve prediction, consistency, and transparency, yet they can also reproduce bias, obscure accountability, and create a false sense of objectivity if they are not evaluated carefully (Kleinberg et al., 2018; Levy, Chasalow and Riley, 2021). Whether algorithmic reform reduces tax inequality is therefore an empirical question.

Cook County provides an unusually informative setting in which to answer it. The county contains more than 2.2 million housing units and 5.1 million residents, making it one of the largest market-based assessment jurisdictions in the United States ¹. It also sits inside a housing market marked by long-running racial and economic segregation, especially across

¹See <https://www.census.gov/quickfacts/fact/table/cookcountyillinois/PST045225>, last accessed on 03/25/2026

Chicago and the south suburbs of Cook County (Novara and Khare, 2017; Acs et al., 2017). Those conditions are precisely the ones under which inexact valuation technology can translate into large distributional distortions. They did: the *Chicago Tribune*/ProPublica investigation ((Grotto, Kambhampati and Long, 2017)) documented systematic overassessment of lower-valued properties under the prior regime, and Berry’s follow-on estimates imply that flawed assessments shifted billions of dollars of property taxes away from the most expensive homes and onto less valuable ones (Berry, 2018).

The reform we study is also unusually consequential. Beginning in 2019 and fully reaching Chicago’s 2021 reassessment, the Cook County Assessor’s Office replaced its legacy township-level linear system with a countywide automated valuation workflow built around gradient-boosted trees, richer geospatial and administrative inputs, a documented production pipeline, and a rare commitment to public release of code, models, and data. We therefore do not study a minor parameter adjustment. We study a bundled modernization of prediction technology, data infrastructure, and institutional transparency. At the same time, Cook County’s triennial “triad” reassessment structure creates useful treatment variation: Chicago adopts the fully modernized residential system in 2021, while the North and South/West suburban triads remain on earlier cycles during the main comparison window (?). This staggered structure allows us to compare treated and untreated areas around the reform while flexibly absorbing reassessment-cycle and COVID-era shocks.

Our empirical design is tailored to the object that matters. The key equity question in assessment is not whether the reform shifts the average level of assessed-to-sale-price ratios, but whether it changes their *gradients*. We therefore estimate an event-study difference-in-differences in *slopes*. The outcome is the log assessment-sale-price ratio, $\log(A/P)$, and the estimands are (i) the *price slope*, the derivative of $\log(A/P)$ with respect to $\log(P)$, and (ii) the *race slope*, the derivative of $\log(A/P)$ with respect to neighborhood minority share. Township-by-quarter fixed effects absorb all local *level* shifts, including reassessment jumps and COVID adjustments, so identification comes from within-township-quarter comparisons of homes that sell at different prices and in neighborhoods with different racial composition. This slope-based DiD is atypical, but it is the right estimand for a reform whose distributional consequences operate through the shape of the assessment schedule rather than its average

level.

A second design choice is equally important: we estimate price and race slopes both separately and jointly. In a segregated housing market, the unconditional race slope combines two conceptually distinct objects. One is a *price channel*: if lower-priced homes are overassessed and minority neighborhoods are lower-priced on average, then price regressivity mechanically generates race regressivity. The other is a *within-price channel*: even holding price fixed, assessment ratios may still vary with neighborhood racial composition because the valuation system pools too coarsely across (larger) geographical areas. Estimating both slopes in the same specification allows us to separate these channels. This decomposition turns out to be essential for interpreting the race results.

We find that the 2021 reform substantially reduced price-based regressivity. Relative to a 2020 baseline price slope of -0.384 , the treatment effect on the price slope is $+0.154$ in the first post-reform quarter and remains positive and statistically significant throughout the next eight quarters, peaking at $+0.223$ by Post3. In economically intuitive terms, for two homes that differ by one log point in price, the gap in assessment ratios falls from about 47% before the reform to about 26% immediately after and to about 17% by Post3. For illustration, compare a \$200,000 home and a \$543,000 home in Chicago. The Post3 compression implies roughly \$60,000 less assessed market value for the lower-priced home than the pre-reform gradient would imply; using Chicago’s estimated 2022 residential effective property tax rate of 1.69%, that corresponds to about \$1,000 per year in tax liability ². These gains are durable: when we repeat the design using the post-appeal certified values, the price-slope effects are nearly identical, indicating that appeals do not materially undo the reform’s main equity improvement.

The race results are more nuanced. Without controlling for price, the post-reform *unconditional* race slope becomes more negative, implying that higher-minority neighborhoods move toward lower assessment ratios on average. But this unconditional pattern is not the cleanest causal object: it exhibits meaningful pre-trend movement and more importantly, in a segregated market, necessarily loads the price channel. Once we jointly estimate price and race slopes, the post-reform coefficients on the *within-price* race slope are modest and

²See https://www.civicfed.org/Effective_Property_Tax_2013_2022, last accessed on 03/25/2026.

mostly positive relative to an already negative baseline conditional race slope. The clearest interpretation is that the reform’s dominant distributional effect is a large flattening of the price gradient, and that much of the observed racial reordering in total assessment ratios operates through that price channel rather than through a large direct shift in within-price treatment by neighborhood race.

The mechanism we emphasize does not require the new system to use race explicitly; instead, it works by reducing two forms of coarse pooling in assessment. Under the legacy system, township-level linear models pooled very different submarkets together, generating a steep negative price gradient in assessment ratios, with lower-priced homes systematically overassessed relative to higher-priced homes. In a segregated housing market, that price error mechanically translates into racial disparity because higher-minority neighborhoods tend to have lower-priced housing stock. A second channel operates within price bins: when the model relies too heavily on coarse neighborhood averages, similarly priced homes can still receive different assessment ratios depending on the neighborhoods with which they are pooled. The 2021 reform reduced both forms of pooling by combining richer fine-geographic information with a more flexible machine-learning model that allows the relationship between location, housing characteristics, and value to vary more finely across space. We show evidence consistent with this interpretation: by 2021, fine geography explains substantially more residual variation in both assessed values and within-price assessment ratios, while the incremental role of coarse neighborhood averages and price itself declines. This helps explain why the reform generated a large compression of price regressivity and a change in the unconditional race gradient, even though the within-price race gradient moved only modestly.

This paper makes several contributions. First, it provides causal evidence on whether a real-world algorithmic modernization in a core government function reduced inequality, rather than merely documenting cross-sectional disparities. Second, it demonstrates why race effects in assessment should be decomposed into price-driven and within-price components in segregated housing markets. Third, by comparing mailed and certified values, it shows that post-assessment appeals, while unequal in incidence, are too limited to reverse the reform’s main upstream improvement in vertical equity.

More broadly, our findings carry practical lessons for governments considering similar rollouts.³ The main equity gains do not appear to come from a race-specific correction or from downstream appeals; they come from improving the initial mapping from property attributes and fine geography to value. That suggests that governments can reduce longstanding fiscal inequities by modernizing data, model flexibility, and transparency at the front end of decision making. But it also suggests a specific governance lesson: agencies should audit *distributional gradients*—not just average predictive accuracy—and should evaluate final outcomes after administrative review, since neither overall fit statistics nor “human oversight” guarantees equitable incidence.

The rest of the paper proceeds as follows. Section 3 describes Cook County’s legacy and post-reform assessment systems. Section 4 introduces the data and model-free evidence. Section 5 presents the slope-based DiD design and main results. Section 5.4 studies appeals and certified values. Section 6 develops and tests a two-pooling framework for understanding racial disparities in assessments.

2 Literature Review

This paper sits at the intersection of three literatures: property-tax inequality, algorithmic distributional effects and bias, and algorithms in the public sector.

Property-tax inequality and assessment regressivity. A substantial literature shows that property assessments often violate the ad valorem ideal. Berry (2021) documents pervasive price regressivity in residential assessments, while Amornsiripanitch (2020) shows that owners of inexpensive houses pay substantially higher effective property tax rates than owners of expensive houses within the same jurisdiction. On the racial dimension, Avenancio-

³Cook County’s experience is not unique: other large assessment jurisdictions are already piloting or deploying related technologies. In 2025, New York City’s Department of Finance awarded a six-month pilot to explore AI-based valuation of residential condominium properties (see <https://a856-cityrecord.nyc.gov/RequestDetail/20250304005>, last accessed on 03/25/2026). Maricopa County, Arizona has already rolled out—and continues to refine—Assessment Analyst GAMA, a cloud-based mass appraisal system that supports more than 1.8 million property assessments annually (see <https://resources.esri.ca/customer-stories/maricopa-county-goes-live-with-assessment-analyst-gama>, last accessed on 03/25/2026).

León and Howard (2022) document a nationwide “assessment gap” under which Black and Hispanic homeowners face higher property tax burdens than white homeowners holding jurisdiction and statutory tax rates fixed. Historical work places these disparities in a longer lineage of racialized tax administration and dispossession (Kahrl, 2024). Relative to this literature, our contribution is not primarily to document that assessment inequality exists, but to estimate whether a concrete institutional reform can causally reduce it.

Recent work begins to unpack why property-tax burdens differ across groups. Avenancio-León and Howard (2022) emphasize that assessments may be less sensitive to neighborhood attributes than market prices are, generating racial disparities in tax burdens in segregated housing markets. Ihlanfeldt and Rodgers (2021) and Ihlanfeldt and Rodgers (2023) show that racial gaps can emerge at multiple stages of the property-tax process, including the initial assessment, exemptions, and taxable value rules. Holz, Novgorodsky and Simon (2024) further show that information frictions and other barriers matter for appeal behavior. Our paper complements this literature in two ways. First, we focus on an institutional change at the initial valuation stage rather than on household take-up or downstream relief programs. Second, by jointly estimating price and race slopes, we provide a decomposition of total race regressivity into a price-driven component and a residual within-price component. That decomposition helps explain why race gradients can move sharply when the price gradient flattens.

Algorithms and fairness: promise and perils This paper contributes to a growing literature on the distributional consequences of algorithmic decision-making, showing that algorithms can sometimes improve both accuracy and equity by replacing noisier or more distorted decision rules (e.g., (Kleinberg et al., 2018; Arnold, Dobbie and Hull, 2025; Zhang et al., 2021; Fu et al., 2025)). We contribute to this literature by providing causal evidence from a core government function—property tax assessment—showing that algorithmic modernization can reduce inequality in tax incidence through a better mapping from property and location characteristics to value.

At the same time, our paper also speaks to the literature on algorithmic bias, which emphasizes that algorithms can create, amplify, or perpetuate disparities when they are

trained on distorted data, optimized for narrow objectives, or deployed in unequal social environments. Prior work shows that even race- or gender-neutral algorithms can generate unequal outcomes in settings such as ad delivery, tax enforcement, and job recommendation ((Lambrech and Tucker, 2019; Elzayn et al., 2025; Zhang and Kuhn, 2024; Bohren, Hull and Imas, 2025)). Our results add to this literature by showing that, in a segregated housing market, racialized burdens can arise through model pooling and valuation error even when race is not used directly; conversely, reducing coarse pooling can mitigate those disparities.

Algorithms in the public sector. More broadly, our paper contributes to the literature on algorithms in government. Existing work emphasizes that algorithmic systems can improve prediction and consistency in public decisions, but also create concerns about opacity, accountability, and social inequity (Kleinberg et al., 2018; Levy, Chasalow and Riley, 2021; Green, 2022). Most of that literature studies settings such as criminal justice, education, or benefits administration. Property assessment is a similarly high-stakes public decision, but it has received much less attention in the algorithms literature despite directly shaping household tax burdens and local public finance. Our results show that the public-sector algorithm debate should pay closer attention to fiscal administration, and that algorithmic systems should be judged not only by mean error but by how they reshape socially salient gradients such as price and race.

3 Institutional Setting

3.1 The Traditional Assessment System (pre-2021)

Prior to the algorithmic transition studied in this paper, Cook County employed a traditional assessment work flow: statistical models estimated on property sales were combined with staff review and a formal appeals process to produce assessed values. The core valuation step used linear hedonic regression models that relate recent transaction prices to observed housing characteristics and location attributes (e.g., size, age, and geography), a standard approach in both real estate economics and assessment practice (Rosen, 1974; Wang and Li, 2019; IIAO, 2011). The estimated coefficients—representing average marginal values such

as price per square foot or value per bedroom—were then applied to unsold properties to generate assessed values.

However, regression models estimated on broad strata (e.g., large geographic groupings or broad property classes) can smooth over neighborhood-specific heterogeneity and help create and/or amplify systematic valuation errors: the model effectively used coefficients that reflected township-wide averages, which were disproportionately influenced by higher-valued properties. When these uniform coefficients were applied to lower-valued properties in economically distressed neighborhoods, the result was systematic overassessment—a key driver of the regressivity documented by the Chicago Tribune investigation (Grotto, Kambhampati and Long, 2017).

Another distinguishing feature of the pre-2021 system was its heavy reliance on undocumented manual adjustments. After the initial model-based valuations were computed, assessors performed manual case-by-case adjustments to individual property assessments (Grotto, Marx and Richards, 2017), without properly maintaining statistics on the frequency or magnitude of these adjustments and declined Freedom of Information Act requests seeking documentation of the hand-check methodology (Grotto, Kambhampati and Long, 2017).

Professional standards emphasize routine sales-ratio studies and diagnostic reporting to evaluate accuracy and equity across property strata (IIAO, 2011). In Cook County, several high-profile independent analyses and media investigations documented substantial price-related bias and dispersion in pre-reform assessed-to-sale ratios (Board, 2018; Berry, 2018). Finally, contemporaneous audits and reporting noted that the office operated with legacy information systems and data constraints, which can limit spatial integration (e.g., GIS) and the ability to iterate on modeling and quality-control tools.⁴

Taken together, the pre-2021 system provides an empirically important baseline: a standard regression-based valuation model that uses limited set of housing characteristics and spatial features, complemented by staff review and appeals, with distributional performance and transparency largely assessed through external studies and investigations.

⁴See <https://www.govtech.com/computing/cook-county-audit-calls-out-old-tech-workforce-shortages.html>, last accessed on 03/25/2026.

3.2 The New Algorithmic Assessment System (2021+)

Beginning in 2019, the Cook County Assessor’s Office (CCAO) rebuilt its residential valuation workflow around an open-source automated valuation model (AVM) and a version-controlled production pipeline. By the 2021 reassessment cycle, first-pass values for single- and multi-family residential parcels were generated using a gradient-boosted decision tree ensemble (LightGBM).

Statistical model. The core methodological change is the use of gradient boosting decision trees, implemented in LightGBM (Friedman, 2001; Ke et al., 2017). Relative to pooled linear hedonic regressions, boosted-tree AVMs can flexibly capture nonlinearities and interactions among characteristics and location without requiring the analyst to pre-specify functional forms. This approach is widely used in modern AVMs in both academic and industry settings, particularly when valuations benefit from high-dimensional covariates and granular spatial information (Ke et al., 2017).

Data, features, and spatial enrichment. CCAO’s published documentation emphasizes that the AVM combines standard parcel and building characteristics with neighborhood and proximity measures (e.g., Census/ACS-based covariates, walkability measures, flood risk proxies, airport noise exposure, and distances to transit), drawing on local administrative and third-party geospatial sources.⁵

Production pipeline and post-modeling review. The CCAO’s AVM is implemented as an end-to-end pipeline (from data ingestion to training, assessment, evaluation, and reporting) with reported model performance using both standard predictive metrics (e.g., RMSE/MAE) and assessor-relevant equity/quality diagnostics (Cook County Assessor’s Office Data Department, N.d.; International Association of Assessing Officers, 2025). The pipeline applies a limited set of explicit post-modeling adjustments and includes a hand-review component before values are finalized (Cook County Assessor’s Office Data Department, N.d.).⁶

⁵For examples of features and sources (including CMAP walkability scores, First Street flood variables, Chicago Department of Aviation noise monitoring, and GTFS feeds for transit stops), see the “Features Used” and “Data Sources” sections in Cook County Assessor’s Office Data Department (N.d.).

⁶For example, CCAO documentation describes post-modeling procedures for multi-card handling, prorated parcels, complex-level averaging for certain housing types, and rounding conventions; see the “Post-Modeling” section of Cook County Assessor’s Office Data Department (N.d.).

Transparency and public documentation. A central institutional change is that the CCAO publicly released its residential modeling code and selected underlying datasets beginning in 2019 and has continued to publish models, technical documentation, and data products through Git hosting and the Cook County Open Data Portal (Cook County Assessor’s Office, 2019; Quaintance, 2019; Cook County Open Data, 2019). The office has also expanded public-facing tools intended to explain assessment drivers using model outputs and local comparable sales.⁷

Implications for identification. For our empirical design, the relevant “treatment” is not only a change in statistical learner, but a bundled modernization of (i) model form (boosted trees), (ii) feature engineering and data/pipeline infrastructure, (iii) documented post-modeling rules and review practices, and (iv) expanded transparency and reporting. Because these components move together within reassessment cycles, our estimates should be interpreted as the combined effect of this package rather than the isolated effect of switching from linear regression to LightGBM. Consistent with this interpretation, an external evaluation of Cook County’s residential assessments during 2019–2024 reports substantial reductions in regressivity relative to earlier years while noting remaining areas for improvement (Berry, 2025).

3.3 Other Important Institutions

Cook County’s Triennial Reassessment Cycle Cook County employs a triennial reassessment cycle, dividing the county into three geographic *triads*: (1) Chicago (8 townships), (2) North suburbs (13 townships), and (3) South & West suburbs (17 townships). Each triad undergoes full reassessment every three years on a staggered schedule—Chicago in years divisible by 3 (e.g., 2015, 2018, 2021), North suburbs one year later (2016, 2019, 2022), and South/West suburbs two years later (2017, 2020, 2023). In non-reassessment years, assessed values carry forward from the most recent reassessment year within each property’s three-year cycle.

⁷For an example of this public communication approach, see CCAO’s release describing a homeowner-facing tool that summarizes the data inputs and sales most influential for an individual valuation (Cook County Assessor’s Office, 2025).

The COVID-19 adjustment (Tax Year 2020) In tax year 2020, the Cook County Assessor’s Office (CCAO) applied a one-time “COVID-19 Adjustment” in response to the anticipated economic and market effects of the COVID-19 pandemic.(Cook County Assessor’s Office, 2020, 2021) For Class 2 residential properties (that our sample fits into), the CCAO used estimated COVID-related changes in unemployment at the Census-tract level—constructed by combining industry-specific unemployment shocks with each tract’s industry employment mix—and mapped this analysis into neighborhood- and property-type-specific percentage reductions in estimated market value (and therefore assessed value).(Cook County Assessor’s Office, 2020) In the South and West suburbs (the 2020 triennial reassessment triad), the resulting reductions for single-family homes and condominiums had a median of 10.3% and ranged from 8.0% to 12.2%.(Cook County Assessor’s Office, 2020) For Chicago and the North suburbs, the CCAO reports reductions for single-family homes and condominiums ranging from 7.5% to 12.3%, with an average of 9.9%.(Cook County Assessor’s Office, 2021)

It is important to note that The COVID adjustment was computed and applied to tax-year 2020 assessed values. Because Cook County assessments are set on a triennial cycle, the resulting 2020 certified values can serve as the baseline carried forward in subsequent non-reassessment years (absent property-specific updates). This means that: (1) For the South triad that was scheduled for full triennial reassessment in 2020, the covid adjustment was also reflected in the 2021 and 2022 assessment. (2) For the North triad that was reassessed in 2019, the covid adjustment would change the 2020 and 2021 assessment values from the 2019 values, until it got reassessed again in 2022. (3) For the Chicago triad, the covid adjustment was only applicable for the year 2020 because it got reassessed in 2021 with the new algorithmic system.

4 Data and Model-free Evidence

4.1 Data and Sample

Our analysis combines several primary data sources from Cook County, Illinois:

1. **Property Assessments** We obtain annual assessed values for all residential proper-

ties from the Cook County Assessor’s Office.⁸ The assessment file includes property identification numbers (PINs), assessment year, property class codes, township and neighborhood codes, and three stages of assessed values: (1) *mailed values*—the assessor’s initial valuation sent to property owners; (2) *certified values*—values after administrative appeals by the Assessor’s Office; and (3) *board values*—final values after the Board of Review’s adjustments. Our primary outcome uses mailed values, as these reflect the direct output of the assessment model before appeals.

2. **Sales Transactions** Property sales data come from the Cook County Assessor’s Office.⁹ The data include PINs, transaction dates, sale prices, and buyer and seller names.
3. **Property Characteristics** We obtain the year-specific home characteristics provided by the CCAO¹⁰, which includes basic home facts (number of bedrooms, full bathrooms, half bathrooms, property type, year built, square footage, lot size) as well as detailed housing characteristics (e.g., garage size, finished attic/basement, wall material, roof material, heating/cooling, etc.)
4. **Appeals** Property Assessment Appeals data are provided by the Cook County Assessor’s Office¹¹, where for each PIN that made an appeal, we observe the mailed assessment value, the certified assessment value, whether the appeal was successful (i.e., appeal lead to a reduced assessed value), and whether the appeal was handled by a professional agent.
5. **Covid Adjustments** We obtain the neighborhood-specific Covid adjustment data provided by the CCAO¹².

⁸See https://datacatalog.cookcountyil.gov/Property-Taxation/Assessor-Assessed-Values/uzyt-m557/about_data, last accessed on 02/27/2026.

⁹See https://datacatalog.cookcountyil.gov/Property-Taxation/Assessor-Parcel-Sales/wvhk-k5uv/about_data, last accessed on 02/27/2026.

¹⁰See https://datacatalog.cookcountyil.gov/Property-Taxation/Assessor-Single-and-Multi-Family-Improvement-x54s-btds/about_data, last accessed on 03/08/2026.

¹¹See https://datacatalog.cookcountyil.gov/Property-Taxation/Assessor-Appeals/y282-6ig3/about_data, last accessed on 02/27/2026.

¹²See https://public.tableau.com/app/profile/ccao/viz/COVID19Dashboard_15906776254570/COVID19AdjustmentMap, last accessed on 03/03/2026.

6. **Demographics** We obtain racial composition (% Black and % Hispanic) for each Census Block, Census Block Group, and Census Tract of Cook County from U.S. Census Bureau and American Community Survey. Properties are linked to these above-mentioned geo classifications via a crosswalk provided by Cook County Assessor’s Office¹³.

Throughout the analysis, we focus on properties that were ever sold in the sample period because our analyses require the sale price as the ground truth to compare the assessed value to. We acknowledge and discuss the endogeneity issues and remedies in Section ???. We choose 2015-2022 because (1) 2015 is the start of the 3-year reassessment cycle of 2015-2017 for the Chicago triad, giving us two full cycles before the focal algorithm adoption in 2021; (2) We stop at 2022 because South triad started to adopt the algorithm in 2023; (3) While we keep 2015-2022 for summary stats and model free evidence, we drop 2022 for North triad because it adopted the algorithm in 2022. For the model-free evidence, we extend the main sample to the end of 2025 for all three triads so that readers can have a more complete and up-to-date view.

We construct our analysis sample through the following steps. We begin with all properties that (1) sold at least once during 2015–2022, (2) are classified as residential single family homes¹⁴, (3) have valid assessed values in the year of sale¹⁵, (4) are inferred as arms-length transactions since prices of non-arms-length sales are not reflective of the fair market value¹⁶, and (5) sold for at least \$10,000¹⁷. In addition, we remove outliers in sale prices above the 99.5th percentile. We construct the base sample at the property-year-quarter level where the observed sale price is property-year-quarter specific and the assessed value is property-cycle(3

¹³See https://datacatalog.cookcountyil.gov/Property-Taxation/Assessor-Parcel-Universe/nj4t-kc8j/about_data, last accessed on 02/27/2026.

¹⁴We focus on various single family home types as classified by the Assessor’s Office: https://prodassets.cookcountyassessoril.gov/s3fs-public/form_documents/Class_codes_definitions_12.16.24.pdf, last accessed on 02/27/2026. Specifically, we use the following property classes: 202, 203, 204, 205, 206, 207, 208, 209, 210, 234, 278, and 295.

¹⁵In the rare event of missing assessed values for the year of sale, we use the assessed value from the previous reassessment cycle before the year of sale, otherwise we drop this sale altogether from the data

¹⁶We remove deed types that are plausibly non-arms-length transactions between willing buyers and sellers, such as quit claim deed, sales between related parties, executor deed, etc. We have removed approximately 8% of the sample that are inferred as non-arms-lengths transactions

¹⁷This is in line with CCAO’s practice of removing outliers in sale prices. Sale prices under \$10,000 can be measurement errors and/or non-arms-length transactions, both of which would affect our analyses.

years) specific.

Our primary outcome variable is the log assessment ratio: $y_{it} = \ln\left(\frac{\text{Mailed Assessed Value}_{it} \times 10}{\text{Sale Price}_{it}}\right)$, where i indexes properties and t indexes time (year-quarter). The factor of 10 converts Cook County’s 10% assessment level to a full market value basis, making y_{it} interpretable as the log ratio of estimated market value to the actual market value. This measure captures proportional assessment error: $y_{it} = 0$ indicates perfect accuracy, $y_{it} > 0$ indicates overassessment, and $y_{it} < 0$ indicates underassessment. We focus on the log ratio rather than the simple ratio because it is less prone to the skewed distribution by effectively symmetrizing assessment errors¹⁸ and more aligned with both the industry practice by assessors¹⁹ and property tax research (e.g., Avenancio-León and Howard (2022)). To prevent extreme values from affecting our results, we winsorize this ratio at the 0.5th and 99.5th percentiles.

Our final analysis sample consists of 227,137 property-sale observations from 2015–2022, spanning 1,247 census tracts and 38 townships. The sample is geographically diverse, with substantial representation from all three triads: Chicago accounts for 45% of observations, North suburbs 28%, and South/West suburbs 27%.


4.2 Model-Free Evidence on Regressivity

Before turning to our difference-in-differences (DiD) design, Figures 1–2 provide model-free evidence on how assessment ratios evolved across the Cook County reassessment cycles and around the 2020 COVID adjustment and the 2021 algorithmic reform. Each point is the mean of $\log(\text{Assessed}/\text{Sale Price})$ for a group of arm’s-length sales in a given triad and year-quarter. Sale-price quintiles (Figure 1) are defined within triad-by-quarter cells, so that properties sold across years within the same triad are comparable. For the race slope plots (Figure 2), we categorize CBGs into quintiles using their minority share.

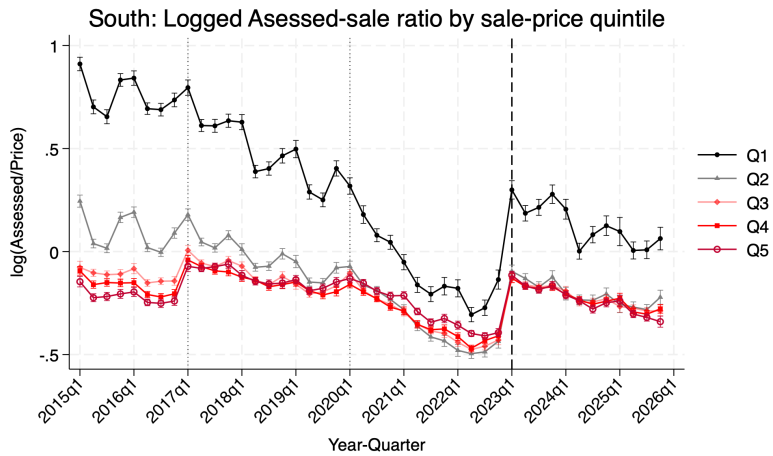
Price-based patterns. Figure 1 shows strong and persistent *price-related* dispersion in assessment ratios under the pre-reform system: within a triad-quarter, low-priced sales (Q1)

¹⁸For example, under simple ratios, a ratio of 10 is drastically further from the mean than its reciprocal of 0.1.


¹⁹The International Association of Assessing Officers (IAAO) explicitly supports the use of log ratios to correct for the non-normal distributions in assesment data. See https://www.iaao.org/wp-content/uploads/Standard_on_Ratio_Studies.pdf, last accessed on 02/17/2026



(a) Chicago Triad



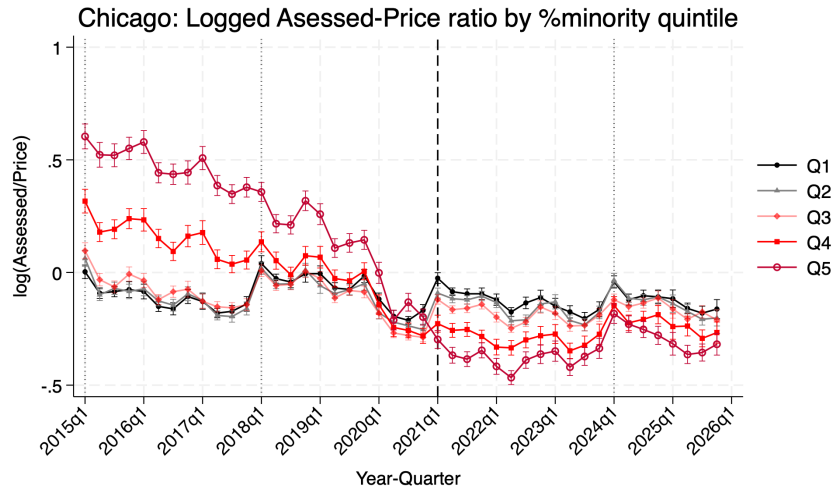
(b) South Triad




(c) North Triad

Figure 1: Model-Free Evidence: Assessment Regressivity by Sale Price Quintile

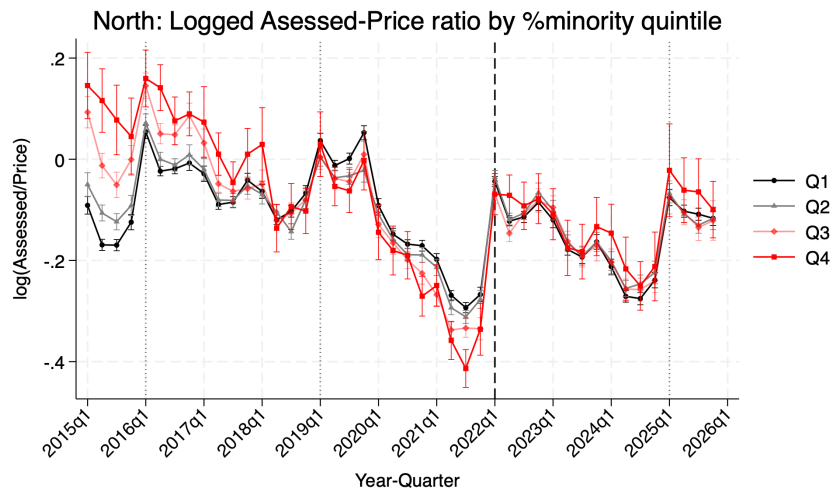
Notes: Each panel plots mean logged(assessed/sale price) by property value quintile (Q1 = lowest 20%, Q5 = highest 20%), where the price quintiles are computed within each triad-quarter. Vertical dotted lines indicate reassessment years for each triad. The black dashed lines (e.g., 2021 for Chicago) marks algorithmic reform adoption.



(a) Chicago Triad



(b) South Triad



(c) North Triad

Figure 2: Model-Free Evidence: Assessment Regressivity by %Minority Quintile

Notes: Each panel plots mean log assessment ratio by census tract % Black quintile (Q1 = lowest 20% Black, Q5 = highest 20% Black). Patterns mirror price-based regressivity because predominantly Black neighborhoods have lower home values. Algorithmic reform reduced racial gaps by reducing price-based regressivity, not through race-conscious adjustments.

tend to have substantially higher log assessment ratios than high-priced sales (Q5). In Chicago, this vertical dispersion is large and fairly stable through 2015–2017, then narrows after the 2018 reassessment, and changes again around 2020–2021. In particular, the lines move sharply during 2020 (concurrent with the covid adjustment), and thereafter the cross-quintile spread in Chicago is markedly compressed relative to earlier years. While the post-2021 relationship is not perfectly flat (the ordering is less monotone than in the early period), the magnitude of the gap between the bottom and top of the price distribution is much smaller than before.²⁰

The suburban triads exhibit the same basic price-gradient—low-priced transactions carry higher ratios than high-priced ones—but the timing of abrupt movements differs. In the South triad, a conspicuous shift occurs around the 2020 reassessment (the triad’s scheduled reassessment year), whereas there is no comparably sharp break in 2021. In the North triad, pronounced level shifts occur at its reassessment years (e.g., 2016, 2019, 2022, 2025), and the largest post-2021 discontinuity aligns with the North triad’s 2022 cycle rather than 2021.²¹

Race-based patterns. Figure 2 groups sales by neighborhood racial composition. In Chicago, the pre-2018 pattern is that higher-minority areas (top quintile) tend to have higher assessment ratios than low-minority areas, with the gap shrinking after 2018. Beginning in 2020 and continuing through 2021–2022, Chicago exhibits a clear *reordering*: the highest-minority quintile moves from being the highest assessment-ratio group to among the lowest, so that the post-2021 pattern looks “race-progressive” in the sense that higher-minority areas have lower assessment ratios on average. The South triad shows a qualitatively similar pre-2019 race regressivity and a similar post-2020 decline for higher-minority quintiles. The total race regressivity seems to be mostly gone starting with the 2023 reassessment, when the South triad started to adopt the new algorithmic system. While the overall pattern is qualitatively similar for the North triad, the North triad displays smaller and noisier racial

²⁰The sharp 2020 movements are visible across multiple series and are not unique to a single quintile, consistent with a countywide shock to market conditions and/or assessment adjustments in that year. Because the figures are purely descriptive, we do not attribute the 2020 shift to a specific mechanism here.

²¹Vertical lines in the figures mark the start of each triad’s reassessment cycle; the thick dashed line marks the first cycle under the post-2021 system for that triad (Chicago in 2021; North in 2022).

gradients.²²

5 Effect on Tax Regressivity: Identification and Results

5.1 Difference-in-Differences Design: Estimating Changes in Slopes

The model-free figures highlight two empirical challenges for identifying the effects of Cook County’s assessment reforms. First, assessment ratios exhibit substantial *level* movements at triad reassessment boundaries and during the COVID period, and these level shifts are not synchronized across space (Figures 1–2). Second, the patterns of interest are not primarily changes in average assessment ratios, but changes in *gradients*: how the log assessment ratio varies with sale price (price regressivity) and with neighborhood racial composition (race regressivity). These considerations motivate a difference-in-differences (DiD) design that isolates *changes in slopes* while absorbing local *level* shocks.

Outcome and gradients. For each arm’s-length sale i in township j and quarter t , we study $y_{ijt} \equiv \log\left(\frac{A_{ijt}}{P_{ijt}}\right)$, where A_{ijt} is the mailed assessed value (our primary administrative output) and P_{ijt} is the transaction price. We define $p_{ijt} \equiv \log(P_{ijt})$ and let m_{ijt} denote neighborhood minority share for the property’s location.²³ A negative relationship between y_{ijt} and p_{ijt} indicates price regressivity (lower-priced sales having higher assessment ratios), while a positive relationship between y_{ijt} and m_{ijt} indicates higher assessment ratios in higher-minority neighborhoods.

Event-study DiD in slopes. Let Treat_j indicate the treated geography (e.g., Chicago townships in the main reform-timing design) and let $\{1[\tau = t - k]\}_k$ denote event-time indicators relative to the reform quarter (normalized so that $k = 0$ is the first post-reform

²²In the North triad, minority-share variation is more limited; the plotted groups therefore show less separation and are estimated with greater uncertainty, especially in some quarters.

²³Our baseline uses CBG-level minority share, but the main results are robust to coarser and finer geographies (Census tract and Census block, respectively), which is consistent with the descriptive evidence that the key variation is not an artifact of tract aggregation.

quarter and $k = -1$ is the omitted reference period).²⁴ The central estimating equation models *changes* in the price and race gradients as:

$$y_{ijt} = \alpha_{jt} + \alpha_{c(i),yr(t)} + (\beta_j + \beta_t) p_{ijt} + (\phi_j + \phi_t) m_{ijt} + \sum_{k \neq -1} \delta_k \text{Treat}_j \cdot 1[\tau = t - k] \cdot p_{ijt} + \sum_{k \neq -1} \rho_k \text{Treat}_j \cdot 1[\tau = t - k] \cdot m_{ijt} + \varepsilon_{ijt}. \quad (1)$$

We include township-by-quarter fixed effects α_{jt} to absorb all local level shocks to assessment ratios (including reassessment-cycle jumps and COVID-era level shifts), so that any slope estimates should be thought of as the regressivity across homes within the township. This choice of within-township regressivity, aligning with other studies (e.g., ?), is reasonable because homes within the same township face the same public amenities and many local tax regulations are made at the township level. We also allow the *baseline* price and race gradients to vary across townships (β_j, ϕ_j) and across calendar quarters (β_t, ϕ_t) , so identification comes from deviations in slope dynamics for treated townships relative to controls, not from assuming common gradients across space or time. Standard errors are clustered at the township level.

The event-study coefficients $\{\delta_k\}$ trace the dynamic effect of the reform on the *price slope* of y_{ijt} (i.e., changes in price regressivity), while $\{\rho_k\}$ trace the dynamic effect on the *race slope* of y_{ijt} (i.e., how the assessment ratio varies with minority share).²⁵

Identification: parallel trends in gradients. Equation (1) is identified by within-township-quarter variation in transaction characteristics: in a given town-quarter, we observe sales at different prices and in neighborhoods with different minority shares. The key identifying assumption is a *parallel trends* condition for slopes: absent the reform, the evolution of the price gradient and race gradient in treated townships would have followed the same path

²⁴We normalize the event study at $k = -1$ (the quarter immediately preceding the reform quarter) so that post- k coefficients are interpretable as changes net of contemporaneous shocks in 2020–2021 (including COVID-era adjustments). Appendix figures re-center the same estimates to a pre-COVID baseline for descriptive purposes.

²⁵In the tables we also report the pre-reform (reference-period) slopes estimated in 2020 as summary rows to aid interpretation, since δ_k and ρ_k are changes relative to the omitted quarter.

as in control townships, after accounting for flexible township-specific and quarter-specific baseline gradients. We assess this assumption using event-study *leads* (pre-period δ_k and ρ_k): the model-free figures suggest that sharp movements in levels occur at reassessment boundaries and in 2020, so the relevant validation is that the treated and control groups do not exhibit differential pre-trends in the *gradients* in the periods immediately preceding the reform.

Why we jointly model price and race slopes. We will start by estimating the effects on the price slope and on the race slope separately. These estimates will inform us the effect on the *total* price regressivity and the *total* race regressivity.

That said, a key observation from the model-free evidence is that price and race patterns are tightly intertwined. Price-based dispersion is large in early years and compresses markedly around the reform period (Figure 1), while the race-gradient in Chicago appears to “flip” in sign around 2020–2021 (Figure 2). Interpreting the race pattern without modeling price can be misleading because neighborhood minority share is correlated with price in a segregated housing market. Specifically, the *total* race regressivity can be decomposed into (i) a price-driven component and (ii) a residual within-price component.

When the reform reduces price regressivity, part (i) mechanically reduces the unconditional race slope even if the residual component changes little. This is consistent with the model-free plots: large changes in the price gradient coincide with a reordering of the unconditional race gradient. As we will show in Section 6, this price-race decomposition provides a coherent basis for a better understanding of the evolution of the race inequality in home value assessment and the role of the algorithmic reform.

5.2 Main results: effects of the algorithmic reform on price and race slopes

Table 1 reports our primary DiD estimates of how Chicago’s 2021 algorithmic reform changed two slopes of the log assessment ratio, $y = \log(A/P)$: (i) the *price slope* with respect to $\log(P)$ (price-based regressivity), and (ii) the *race slope* with respect to neighborhood minority share (race-based regressivity). All specifications absorb township-by-quarter fixed

Table 1: Main DiD: Price Regressivity and Racial Bias

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.032 (0.024)		-0.021 (0.031)
Treat × Pre3 × ln(price)	0.035 (0.026)		0.028 (0.036)
Treat × Pre2 × ln(price)	0.014 (0.014)		0.013 (0.018)
Treat × Post0 × ln(price)	0.154*** (0.042)		0.177*** (0.038)
Treat × Post1 × ln(price)	0.170*** (0.027)		0.175*** (0.027)
Treat × Post2 × ln(price)	0.197*** (0.026)		0.202*** (0.029)
Treat × Post3 × ln(price)	0.223*** (0.038)		0.249*** (0.038)
Treat × Post4 × ln(price)	0.158*** (0.044)		0.178*** (0.044)
Treat × Post5 × ln(price)	0.175*** (0.033)		0.192*** (0.032)
Treat × Post6 × ln(price)	0.139*** (0.040)		0.155*** (0.039)
Treat × Post7 × ln(price)	0.158*** (0.046)		0.188*** (0.048)
Treat × Pre4 × minority		0.075** (0.034)	-0.037 (0.067)
Treat × Pre3 × minority		-0.069 (0.041)	-0.026 (0.045)
Treat × Pre2 × minority		0.009 (0.032)	0.008 (0.033)
Treat × Post0 × minority		-0.149** (0.065)	0.057 (0.037)
Treat × Post1 × minority		-0.116 (0.071)	0.054 (0.041)
Treat × Post2 × minority		-0.102* (0.060)	0.065* (0.037)
Treat × Post3 × minority		-0.124* (0.061)	0.084** (0.039)
Treat × Post4 × minority		-0.119 (0.071)	0.062 (0.060)
Treat × Post5 × minority		-0.106 (0.079)	0.051 (0.045)
Treat × Post6 × minority		-0.098 (0.079)	0.033 (0.040)
Treat × Post7 × minority		-0.059 (0.085)	0.097** (0.042)
2020 price slope	-0.384*** (0.043)		-0.452*** (0.049)
2020 race slope		-0.003 (0.025)	-0.433*** (0.069)
Observations	284,935	284,935	284,935
R-squared	0.604	0.237	0.647
Town-by-quarter FE	YES	YES	YES

Notes: This table presents the main DiD (Equation 1) estimates examining how the algorithmic reform affected how the assessment values vary with logged sale price (the price slope) and Census-Block-Group-level share of minority population (the race slope). Across specifications, the outcome variable is $\ln(\text{Assessed}/\text{Sale Price})$, where Assessed is the initial assessed values. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space but they can be seen in event study plots (see Figure 3 and Figure 4). All specifications include township-by-quarter fixed effects. Standard errors are clustered at the township level and in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

effects, so identification comes from within-township changes in these slopes over event time relative to the reference quarter $t = -1$ (2020Q4). For interpretation, the table also reports the 2020 baseline slope in each specification.

Price regressivity falls sharply and persistently. Column (1) shows a large, immediate, and precisely estimated flattening of the price gradient after the reform. The three pre-period lead coefficients are small and statistically insignificant, while every post-period coefficient on $\text{Treat} \times \text{Post}_k \times \log(P)$ is positive and statistically significant. Relative to the 2020 baseline price slope of -0.384 , the first post-reform estimate of $+0.154$ implies about a 40% reduction in price regressivity, and the effect reaches $+0.223$ by Post3, or about 58% of the baseline slope. In level terms, for two homes that differ by one log point in price, the implied gap in A/P falls from $\exp(0.384) - 1 \approx 47\%$ before the reform to $\exp(0.384 - 0.154) - 1 \approx 26\%$ in Post0 and to $\exp(0.384 - 0.223) - 1 \approx 17\%$ in Post3.

Column (3), which jointly estimates price and race slopes, yields very similar price effects (if anything slightly larger), indicating that the improvement in price-based regressivity is robust to conditioning on neighborhood racial composition. Using the joint specification, the baseline price slope is -0.452 , and the post-reform coefficients range from $+0.155$ to $+0.249$, again implying a substantial and persistent compression of the price gradient.

Race slope: unconditional shifts versus within-price shifts. The race results are more specification-dependent. Column (2) estimates the *unconditional* race gradient, without controlling for $\log(P)$. Here the 2020 baseline race slope is essentially zero (-0.003), but the post-reform coefficients are negative in every quarter, ranging from -0.149 at Post0 to -0.059 at Post7. Read literally, this implies that the overall race slope becomes more negative after the reform. At the same time, this specification is less clean than the price specification: one lead coefficient (Pre4) is positive and statistically significant, and the event-study plot shows substantial movement in the unconditional race slope well before 2021. We therefore interpret Column (2) more cautiously as descriptive evidence on the total race gradient.

Column (3) instead estimates the *within-price* race gradient. In this joint specification, the post-period coefficients are modest and mostly positive, with statistical significance only in a subset of post periods. Because the 2020 conditional race slope is already negative (-0.433), these positive post coefficients imply a partial attenuation of the negative within-price race gradient rather than a reversal of it. Even the largest post estimate, $+0.097$ in

Post7, offsets only about 22% of the baseline conditional race slope, so the within-price race slope remains negative throughout the post period. Put differently, once price is held fixed, the reform does not appear to generate a large direct reweighting by neighborhood minority share.

Event-study evidence. Figures 3–4 reinforce this interpretation. First, the price-slope event studies show a clear break at the reform date. In both the univariate specification (Figure 3a) and the joint specification (Figure 4a), the estimated effect jumps upward at $t = 0$ and remains positive throughout the next eight post quarters, peaking roughly three to four quarters after implementation. The recent pre-period coefficients are close to zero, which supports the identifying assumption for the price results.


Second, the race-slope event studies are more mixed. Without controlling for price (Figure 3b), the unconditional race slope drifts downward well before the reform and remains negative afterward. That pattern cautions against interpreting the unconditional post-2021 race coefficients as a clean break caused solely by the reform. In the joint specification (Figure 4b), the near-treatment pre-period coefficients are closer to zero and the post-treatment coefficients are modestly positive, consistent with Column (3)’s finding of a small attenuation in the negative within-price race slope. Taken together, the dynamic evidence points to a clear conclusion: the algorithmic reform’s most robust distributional effect is a large and persistent reduction in price-based regressivity. Any post-reform racial reordering in the overall assessment ratio appears to operate primarily through that price channel, rather than through a large direct change in within-price racial disparities.

5.3 Robustness Checks

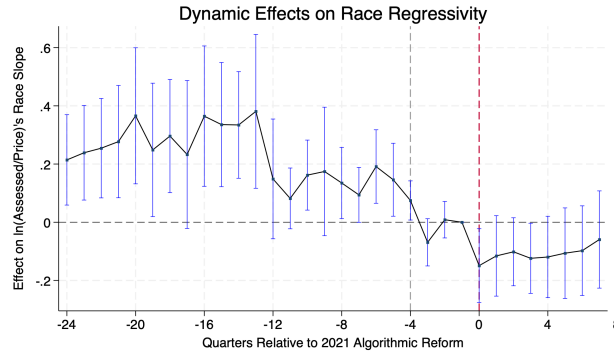
We conduct several robustness checks to address three main concerns: the triad-level rollout of treatment, the measurement of racial composition, and the possibility that highly local housing-market trends confound the main estimates.

1. **Triad-level treatment assignment.** The reform was adopted at the triad-by-year level, so all properties within a triad share the same treatment status. This raises

Figure 3: Dynamic Treatment Effects: Price and Race Regressivity



(a) Treatment Effect on Price Regressivity (Table 1 Column (1))




(b) Treatment Effect on Race Regressivity (Table 1 Column (2))

Notes: Each point represents the estimated coefficient θ_τ from equation (??). Panel (a) interacts treatment with log sale price; Panel (b) interacts treatment with tract-level % Black. Event time $\tau = 0$ corresponds to Chicago's Q1 2021 reassessment. Shaded areas represent 95% confidence intervals clustered at the township level. The vertical dashed line marks the treatment date.


an inference concern because the effective level of variation is much coarser than the parcel-level sample, and with only three triads, results could be sensitive to the choice of control group.

We address this in two ways. First, we re-estimate the baseline DiD while dropping one control triad at a time. The main results remain similar when excluding the North triad (Table A1) or the South triad (Table A2), suggesting that they are not driven by any single comparison group. Second, we report *triad-level wild cluster bootstrapped*

Figure 4: Dynamic Treatment Effects: Price and Race Regressivity Jointly Estimated



(a) Treatment Effect on Price-Based Regressivity
(Table 1 Column (3))



(b) Treatment Effect on Race-Based Regressivity
(Table 1 Column (3))

Notes: Each point represents the estimated coefficient θ_τ from equation (??). Panel (a) interacts treatment with log sale price; Panel (b) interacts treatment with tract-level % Black. Event time $\tau = 0$ corresponds to Chicago's Q1 2021 reassessment. Shaded areas represent 95% confidence intervals clustered at the township level. The vertical dashed line marks the treatment date.

standard errors (Table A3), which provide a more conservative inferential check under very coarse treatment assignment. The post-reform coefficients are generally positive, but with triad-level wild-bootstrap inference and only three clusters, statistical precision is limited; only one post-period estimate is individually significant at conventional levels.

2. **Alternative minority measures.** Our main measure of racial composition is Census Block Group minority share, which matches our focus on neighborhood-level differences

in assessment. To show that the findings do not depend on this specific definition, we replace it with Census Block minority share, a more spatially granular measure. The results remain similar (Table A4).

We also use inferred homeowner race based on the seller’s name (Table A5). This check is useful because it asks whether the results extend from neighborhood racial composition to an owner-level proxy. We find that the price-regressivity results largely persist, while the race-slope results differ somewhat. We do not view this as inconsistent with the main analysis: name-based race imputation is noisy, and owner race is a different estimand from neighborhood racial composition.

3. More local housing-market trends. A final concern is that the main estimates may partly reflect neighborhood-level price dynamics rather than the reform itself, especially because the COVID adjustment was implemented at the neighborhood level. To address this, we replace township fixed effects and baseline slopes with neighborhood fixed effects and neighborhood-specific slopes. As shown in Table A6, the price-regressivity results remain similar, while the race-slope estimates become statistically insignificant, likely because there is limited within-neighborhood variation in CBG minority share.

We also control for Census-tract gentrification interacted with time indicators. This directly addresses the concern that faster appreciation in lower-priced or higher-minority areas could mechanically affect assessment-price ratios. The main results remain unchanged (Table A7).

5.4 Do Appeals Undo the Algorithm’s Equity Improvements?

Our main analysis focuses on the assessor’s *initial* output—the mailed assessed value. In Cook County, however, mailed values can be revised through the administrative appeals process, so the economically relevant object for tax liability is the *certified* value after appeals. This raises a natural concern: if appeals are disproportionately used (and successfully used) by higher-value or lower-minority areas, the appeals stage could partially (or fully) undo the equity gains generated by the algorithmic reform at the mailing stage.

We address this concern in two steps. First, we document *who appeals* and what they

obtain. Second, we re-estimate our main DiD design using the post-appeal certified value as the outcome. The short answer is that appeals are indeed socioeconomically patterned, but they do *not* materially attenuate the reform’s improvement in the price gradient, and the post-appeal results are extremely similar to the mailed-value results.

Who appeals? Table 2 characterizes the appeals using CCAO administrative appeals records. Note that this data is only meaningfully available beginning in 2021²⁶, so Table 2 describes post-reform appeals behavior across all three triennial cohorts. Two patterns are especially salient.

Table 2: Appeals Analysis: Propensity, Outcomes, and Representation

	(1) Appealed (0/1) (All)	(2) Pct Reduction (×100) (All)	(3) Success (0/1) (Appeal=1)	(4) Pct Reduction (×100) (Appeal=1)	(5) Has Agent (0/1) (Appeal=1)
ln(Sale Price)	0.055*** (0.014)	0.247*** (0.068)	0.166*** (0.014)	1.606*** (0.163)	0.016 (0.025)
minority	-0.043*** (0.009)	0.018 (0.022)	0.094*** (0.011)	1.041*** (0.128)	-0.124*** (0.042)
DV Sample Mean	0.1113	0.2019	0.2073	1.7190	0.6185
Observations	4,480,394	4,480,394	498,570	498,570	498,570
R-squared	0.129	0.043	0.059	0.065	0.107
Township × Year FE	Yes	Yes	Yes	Yes	Yes
Class × Year FE	Yes	Yes	Yes	Yes	Yes

Notes: This table presents regressions examining the determinants of appeal behavior and outcomes. Column (1) regresses an indicator for whether the property filed an appeal on log mailed assessment and neighborhood percent Black. Column (2) regresses percentage reduction in assessment (defined as $\frac{\text{Mailed}-\text{Certified}}{\text{Mailed}} \times 100$, which equals zero for non-appellants) on the same predictors. Columns (3)-(5) restrict to properties that appealed. Column (3) examines appeal success (any reduction obtained). Column (4) examines percentage reduction among appellants. Column (5) examines professional representation (agent hired). All specifications include township-by-year and property-class-by-year fixed effects. Standard errors clustered at township level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

First, appeals are strongly increasing in property value. A one log-point increase in sale price is associated with a 5.5 percentage point increase in the probability of filing an appeal (Column 1), which is large relative to a baseline appeal rate of 11.1%. Conditional on appealing, higher-priced properties are also substantially more likely to obtain a reduction (Column 3: +16.6 pp on a baseline success rate of 20.7%) and to obtain larger percentage

²⁶Pre-2021 appeal data are largely missing from the CCAO appeals records.

reductions (Column 4: +1.61 percentage points on a baseline conditional reduction of 1.72%). Thus, the appeals stage mechanically shifts certified values downward more for high-priced properties (Column 2), both because they appeal more and because their appeals are more successful and larger when successful.

Second, appeals behavior varies with neighborhood racial composition in a way consistent with barriers to access on the extensive margin but meaningful relief conditional on participation. Properties in higher-minority neighborhoods are less likely to appeal (Column 1: -0.043 for a unit increase in minority share; roughly -0.43 pp for a 10 percentage-point increase), and among appellants they are less likely to use professional representation (Column 5: -0.124 ; roughly -1.24 pp for a 10 percentage-point increase). At the same time, conditional on appealing, higher-minority neighborhoods have *higher* success rates and larger conditional reductions (Columns 3–4), suggesting selection into appeals: households who do appeal in these neighborhoods may do so when overassessment is particularly salient.²⁷

Taken together, Table 2 confirms that appeals are socioeconomically patterned in ways that could, in principle, reintroduce regressivity through the post-assessment process.

Do appeals undo the reform’s equity effects? Certified-value DiD. To test whether appeals offset the reform’s equity improvements, we replicate our main DiD/event-study design using $y_{ijt}^{\text{cert}} = \log\left(\frac{\text{Certified}_{ijt}}{P_{ijt}}\right)$ as the outcome, i.e., the log assessment ratio computed using certified (post-appeal) values. Table 3 reports the results.

The central finding is that the reform’s effect on the price slope is virtually unchanged when we move from mailed to certified values. In Column (1), the post-period coefficients on $\text{Treat} \times \text{Post}_k \times \ln(P)$ remain positive, large, and precisely estimated in every quarter. Their magnitudes are nearly identical to the mailed-value estimates (for example, Post0 is 0.155 here versus 0.154 in the mailed-value specification, and Post3 is 0.223 in both tables), indicating that the flattening of the price gradient survives the appeals stage.²⁸

Economically, the certified-value results imply a large and persistent reduction in price-based regressivity in the values that ultimately determine tax liability. The baseline 2020

²⁷Because we do not observe counterfactual overassessment among non-appellants, we interpret the conditional patterns descriptively: they are consistent with selection, but not uniquely diagnostic of it.

²⁸If appeals were systematically undoing the reform’s equity gains, the certified-value post coefficients would be materially smaller than the mailed-value coefficients. We do not observe that pattern.

Table 3: Certified: Price Regressivity and Racial Bias

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.026 (0.027)		-0.014 (0.034)
Treat × Pre3 × ln(price)	0.034 (0.027)		0.028 (0.037)
Treat × Pre2 × ln(price)	0.013 (0.015)		0.013 (0.020)
Treat × Post0 × ln(price)	0.155*** (0.041)		0.180*** (0.037)
Treat × Post1 × ln(price)	0.170*** (0.027)		0.176*** (0.026)
Treat × Post2 × ln(price)	0.195*** (0.027)		0.202*** (0.029)
Treat × Post3 × ln(price)	0.223*** (0.039)		0.251*** (0.039)
Treat × Post4 × ln(price)	0.160*** (0.042)		0.180*** (0.043)
Treat × Post5 × ln(price)	0.177*** (0.033)		0.195*** (0.031)
Treat × Post6 × ln(price)	0.139*** (0.040)		0.157*** (0.039)
Treat × Post7 × ln(price)	0.163*** (0.046)		0.193*** (0.048)
Treat × Pre4 × minority		0.069* (0.038)	-0.035 (0.067)
Treat × Pre3 × minority		-0.066 (0.041)	-0.023 (0.047)
Treat × Pre2 × minority		0.012 (0.036)	0.012 (0.037)
Treat × Post0 × minority		-0.145** (0.064)	0.065* (0.037)
Treat × Post1 × minority		-0.110 (0.072)	0.061 (0.041)
Treat × Post2 × minority		-0.090 (0.062)	0.076** (0.037)
Treat × Post3 × minority		-0.119* (0.062)	0.090** (0.040)
Treat × Post4 × minority		-0.125 (0.075)	0.059 (0.065)
Treat × Post5 × minority		-0.105 (0.081)	0.057 (0.046)
Treat × Post6 × minority		-0.086 (0.079)	0.046 (0.040)
Treat × Post7 × minority		-0.057 (0.085)	0.103** (0.042)
2020 price slope	-0.386*** (0.043)		-0.454*** (0.049)
2020 race slope		-0.001 (0.025)	-0.434*** (0.070)
Observations	284,935	284,935	284,935
R-squared	0.597	0.231	0.639
Town-by-quarter FE	YES	YES	YES

Notes: This table presents difference-in-differences (Equation 1) event study estimates examining how the algorithmic reform affected how the post-appeal assessment outcomes vary with logged sale price (the price slope) and Census-Block-Group-level share of minority population (the race slope). Across specifications, the outcome variable is $\ln(\text{Certified}/\text{Sale Price})$, where Certified is the final assessed value after appeals. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space. All specifications include township-by-quarter fixed effects. Standard errors are clustered at the township level and in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

price slope in Column (1) is -0.386 . At Post0, the estimated change is $+0.155$, implying about a 40% reduction in the magnitude of price regressivity ($0.155/0.386 \approx 0.40$). By Post3, the estimated change reaches $+0.223$, or about 58% of the baseline slope. In level terms, for two homes that differ by one log point in price, the implied gap in A/P falls from

$\exp(0.386) - 1 \approx 47\%$ before the reform to $\exp(0.386 - 0.155) - 1 \approx 26\%$ at Post0 and to $\exp(0.386 - 0.223) - 1 \approx 18\%$ at Post3. Column (3), which estimates price and race slopes jointly, yields the same conclusion: the certified-value price effects are, if anything, slightly larger, with post-period coefficients ranging from +0.157 to +0.251 against a baseline slope of -0.454 .

For the race slope, the certified-value results again mirror the mailed-value results in showing that the conclusions depend on whether price is modeled directly. Without controlling for price (Column (2)), the post-period race-slope estimates are negative in every quarter, with statistically significant declines at Post0 and Post3. At the same time, this unconditional specification should be interpreted cautiously: one pre-period lead (Pre4) is also statistically significant, and the baseline 2020 race slope is essentially zero (-0.001). Once price is included (Column (3)), the post-period coefficients become modestly positive, with statistical significance in several quarters (Post0, Post2, Post3, and Post7). Relative to the baseline conditional race slope of -0.434 , however, these are still fairly small movements: even the largest post estimate, +0.103 in Post7, offsets less than one quarter of the baseline conditional race gradient. As in the mailed-value results, the reform's most robust distributional effect is the compression of the price gradient, while changes in the within-price race gradient are more modest and specification-sensitive.

Why appeals do not overturn the slope improvements. Taken together, Tables 2 and 3 suggest that appeals do not materially undo the reform's main equity gains. Although appeals are unequally used and can matter for individual parcels, the countywide gradients that summarize regressivity change very little between mailed and certified values. In particular, the certified-value price coefficients are nearly identical to the mailed-value coefficients throughout the post period, which rules out any large re-steepening of the price gradient after appeals. The certified race results also preserve the same qualitative pattern as in the mailed-value specifications: a more negative unconditional race slope, but only modest positive shifts in the within-price race slope. The broad implication is that the reform's primary equity effect occurs at the initial assessment stage, and the appeals process is not large enough to reverse that improvement in the final certified values.

Overall, the certified-value estimates indicate that Cook County’s algorithmic reform delivers a durable reduction in price regressivity in the values that ultimately determine tax liability. Whatever redistribution occurs during appeals, it does not undo the reform’s central slope-based improvement.

6 Mechanisms: A Two-Pooling Framework for Racial Disparities

Our main DiD results show two facts that, taken together, require a mechanism: (i) the 2021 reform produces a large and persistent flattening of the *price slope* in $\log(A/P)$, while (ii) the *unconditional* race slope can move substantially even when the *within-price* (“residual”) race slope moves little once price is controlled for. This section provides a framework that reconciles these patterns and clarifies why technology can change racial inequality even when race is not an explicit model input.

6.1 Intuition: decomposing the race slope into two pooling channels

Essentially, the *unconditional* race slope θ obeys the following omitted-variable decomposition:

$$\underbrace{\theta}_{\text{unconditional race slope}} = \underbrace{\beta_m}_{\text{within-price race slope}} + \underbrace{\beta_p}_{\text{price slope}} \cdot \underbrace{\pi}_{\text{price-race gradient}}, \quad (2)$$

where β_p captures *price regressivity* (a more negative β_p means lower-priced homes have higher assessment ratios) and β_m captures the *residual race gradient holding price fixed*. Think of π as the slope from $p = \pi m + v$. Equation (2) highlights two distinct channels through which “pooling” in the assessment mapping can generate (or eliminate) racial disparities:

Pooling channel 1: cross-price / cross-neighborhood pooling (the price component). If the valuation process effectively pools heterogeneous markets—e.g., by applying

an overly rigid mapping from characteristics to value across locations or price ranges—then errors in y become systematically related to p (a steep negative β_p). This is likely true with legacy assessment models that utilize linear regression models. Specifically, the cook county pre-reform model uses township-specific linear models that apply the same coefficients of home features to all homes within the township, likely leading to overassessing low-valued properties and underassessing high-valued ones. In a segregated housing market, p and m are strongly negatively related ($\pi < 0$). Thus, even if the assessment system were race-neutral conditional on price ($\beta_m \approx 0$), a steep price slope mechanically induces an unconditional race slope $\theta \approx \beta_p \pi$.

Pooling channel 2: within-neighborhood pooling (the residual race component).

Even after holding price fixed, racial composition may still predict assessment ratios if the system pools too coarsely *within* markets. For example, if the location component is largely an average neighborhood intercept (or if the functional form is insufficiently flexible to capture within-neighborhood heterogeneity), then two homes with similar prices and observable characteristics can nevertheless receive systematically different assessment ratios depending on which neighborhoods they pool with. For example, the same-price minority-owned home might be assessed less than the white counterpart because it is pooled with a lower-valued neighborhood.

In fact, with legacy assessment models, there may exist a *natural tension* between the price slope (β_p) and the within-price race slope (β_m): the (inadvertant) extensive reliance on coarse neighborhood intercepts can help reduce price regressivity (as compared to a pure basic linear model), but it will likely amplify the within-neighborhood pooling, counterbalancing the price regressivity gain.

Why the algorithmic reform can move the unconditional race slope. The reform combines (i) richer *fine geography* information and (ii) more flexible machine learning that allows interactions between location and home characteristics. Both changes reduce the need to rely on coarse geographic averages and allow the mapping from home characteristics and location to adapt across segments of the market. In terms of (2), these technological changes

are expected to (a) compress the magnitude of β_p (less price regressivity) by reducing cross-market pooling, and (b) potentially change β_m by reducing within-neighborhood pooling. Because θ is the sum of these two components, a large reduction in the price component $\beta_p\pi$ can flip the sign of the unconditional race slope even if the residual component β_m changes little or only modestly. This logic motivates our emphasis on jointly estimating price and race slopes in the DiD: it separates movements in β_p from movements in β_m .

6.2 Evidence: unpooling from coarse neighborhood averages toward fine geography

In this section, we test whether assessed values and assessment ratios increasingly reflect *fine* geography (micro-location) rather than coarse neighborhood averages. We do so using partial- R^2 decompositions that compare the incremental explanatory power of (i) coarse neighborhood fixed effects versus (ii) fine geography controls, holding a rich set of housing characteristics fixed.²⁹

Table 4 summarizes the results for three cross-sections in Chicago (2015, 2018, 2021), using (Panel A) all assessed parcels with outcome $\log(A)$ and (Panel B) the sales sample with outcome $\log(A/P)$ and flexible within-year price bins.

Overall unpooling in assessed values (Panel A). Panel A shows that, in the earlier cross-section (2015), coarse neighborhood fixed effects explain substantial residual variation in $\log(A)$ even after controlling for housing structure and micro-location variables (partial $R^2 = 0.1449$), whereas fine geography explains relatively little residual variation once neighborhood effects are included (partial $R^2 = 0.0377$). By 2021, this relationship reverses: fine geography explains a large share of residual variation beyond neighborhood averages (partial $R^2 = 0.2434$) while the incremental contribution of neighborhood fixed effects conditional on fine geography is smaller (partial $R^2 = 0.1248$). This pattern is consistent with a shift away

²⁹For each year, we estimate a “full” regression and a “reduced” regression on the same sample and compute partial $R^2 \equiv 1 - \text{RSS}_{\text{full}}/\text{RSS}_{\text{reduced}}$. A higher partial R^2 for a block of variables indicates that block explains more residual variation net of the controls held fixed. The fine-geography block includes school-zone fixed effects (constructed from elementary and secondary school identifiers) and micro-location variables (e.g., flood risk and airport noise exposure). Neighborhood fixed effects capture coarse location averages.

Table 4: Evidence for the Two-Pooling Mechanism: Partial R^2 Decompositions (Chicago)

	2015	2018	2021
Panel A.			
Overall assessment levels (all homes): outcome $\log(A)$			
Coarse neighborhood FE X + fine geography	0.1449	0.1864	0.1248
Fine geography (school + micro-geo) X + neighborhood FE	0.0377	0.1792	0.2434
Panel B.			
Within-price assessment ratios (sales only): outcome $\log(A/P)$			
Coarse neighborhood FE X + price bins + fine geography	0.1231	0.1494	0.1163
Fine geography (school + micro-geo) X + price bins + neighborhood FE	0.0907	0.1628	0.2580
Price bins (100) X + neighborhood FE + fine geography	0.8587	0.8670	0.7269

Notes: X denotes a rich set of housing characteristics (beds/baths, year built, square footage and lot size, basement/attic/garage indicators, exterior materials, heating/cooling, design plan, property class indicators). “Fine geography” includes school-zone fixed effects and micro-location variables (flood exposure and airport-noise measures). Panel B price bins are 100 within-year quantiles of $\log(P)$, absorbing an extremely flexible nonparametric relationship between assessment ratios and sale price.

from coarse neighborhood averaging toward more granular within-neighborhood valuation—precisely what one would expect from a system that incorporates richer GIS-linked data and flexible location-by-structure interactions.

Within-price unpooling from neighborhood mean in assessment ratios (Panel B).

Panel B connects more directly to our regressivity estimands. Because the full model absorbs 100 within-year price bins, the remaining variation is, by construction, *within* narrow price ranges. In this setting, neighborhood fixed effects capture how much assessment ratios still differ across neighborhoods for similarly priced homes (a within-price pooling diagnostic), while fine geography captures whether micro-location can explain those within-price differences.

Two results stand out. First, the incremental explanatory power of fine geography conditional on neighborhood FE rises sharply by 2021 (from 0.0907 in 2015 to 0.2580 in 2021), indicating that post-reform assessment ratios vary systematically with micro-location even within neighborhoods and within price ranges. Second, the incremental contribution of coarse neighborhood FE conditional on fine geography is smaller in 2021 than earlier years (0.1163 in 2021 versus 0.1231 in 2015 and 0.1494 in 2018), consistent with less reliance on

coarse neighborhood averages once fine geography is available.³⁰

Finally, the partial R^2 of price bins (last row) declines markedly in 2021 (0.7269 vs. about 0.86 in earlier years). Because price bins absorb a flexible price relationship, this drop indicates that less of the remaining variation in $\log(A/P)$ is structured by price once housing characteristics and location are controlled for—consistent with our main DiD finding that the reform substantially flattens the price slope (reduced price regressivity).

Implications for racial inequality. Together, the partial- R^2 evidence supports the two-pooling interpretation of the main results. The rise of fine geography in explaining both assessment levels and within-price assessment ratios is consistent with an algorithmic system that uses granular spatial data and flexible machine learning to reduce cross-market pooling (compressing β_p) and to better capture within-market heterogeneity (affecting β_m). Through the decomposition in (2), these changes provide a coherent explanation for why the unconditional race slope can shift substantially when the price slope flattens, even if the within-price race gradient changes more modestly.

7 Concluding Remarks

This paper studies whether an algorithmic reform can reduce inequality in the residential property assessment. We examine Cook County’s 2021 reform, which replaced a legacy assessment workflow with a more flexible automated valuation system built on richer housing and geospatial data, documented production rules, and greater public transparency. Using the county’s staggered triennial reassessment structure, we find that the reform produced a large and persistent reduction in *price regressivity* and modestly improved the within-price race slope by making it more positive. These two changes together makes the post-reform unconditional race slope more negative, suggesting that higher-minority neighborhoods move toward lower assessment ratios on average.

The findings carry practical implications for assessment offices and other public agencies

³⁰Panel B should not be interpreted as implying neighborhood effects become irrelevant; rather, it shows that the relative importance of fine geography increases substantially post-reform in explaining within-price dispersion in assessment ratios.

considering similar reforms. First, the largest equity gains appear to come from improving the initial administrative decision, not from relying on downstream correction through appeals. Second, richer and more granular data matter: part of the benefit of algorithmic reform is not merely a new learner, but a broader modernization of data, geography, pipeline infrastructure, and transparency. Third, fairness evaluation should distinguish between total disparities and disparities conditional on economically relevant covariates. In our setting, that means separating the total race gradient from the within-price race gradient; analogous decompositions may be equally important in other public-sector applications.

Several limitations suggest directions for future research. Our estimates capture the effect of a bundled institutional reform rather than the isolated contribution of any one component, such as switching from linear regression to gradient boosting. Our sales-based design focuses on properties that transact and therefore uses sale price as the benchmark for market value; future work could explore complementary designs for the broader parcel universe. It would also be valuable to trace the reform further downstream, including Board of Review outcomes, exemption take-up, and eventual tax payments, and to study later suburban rollouts and comparable reforms in other jurisdictions. More generally, an important next step is to understand when algorithmic modernization complements administrative capacity and when it may instead reproduce old inequities in a new technical language.

Taken together, the results suggest a cautiously optimistic conclusion. In Cook County, algorithmic reform did not simply make assessments more modern; it made them meaningfully less regressive with respect to price, and it did so in the values that ultimately determine tax liability. That does not eliminate all inequities in the property-tax system. But it does show that long-standing distributional distortions in public administration can be reduced when governments invest in better data, more flexible models, and transparent evaluation of who benefits and who bears the burden.

References

- Acs, Gregory, Rolf Pendall, Mark Treskon and Amy Khare. 2017. The cost of segregation: National trends and the case of Chicago, 1990–2010. Technical report Urban Institute.
- Amornsiripanitch, Natee. 2020. “Why are residential property tax rates regressive?” *Available at SSRN 3729072* .
- Arnold, David, Will Dobbie and Peter Hull. 2025. “Building nondiscriminatory algorithms in selected data.” *American Economic Review: Insights* 7(2):231–249.
- Atuahene, Bernadette and Christopher Berry. 2018. “Taxed out: Illegal property tax assessments and the epidemic of tax foreclosures in Detroit.” *UC Irvine L. Rev.* 9:847.
- Avenancio-León, Carlos F and Troup Howard. 2022. “The assessment gap: Racial inequalities in property taxation.” *The Quarterly Journal of Economics* 137(3):1383–1434.
- Berry, Christopher. 2018. “Estimating Property Tax Shifting Due to Regressive Assessments: An Analysis of Chicago, 2011 to 2015.” *Center for Municipal Finance, The University of Chicago*. <https://propertytaxproject.uchicago.edu/related-research-2-3> .
- Berry, Christopher. 2025. “An Evaluation of Progress on Residential Assessment Fairness in Cook County.” https://bpb-us-w2.wpmucdn.com/voices.uchicago.edu/dist/6/2330/files/2025/09/Kaegi-Evaluation-9_04.pdf. Accessed 2026-02-27.
- Berry, Christopher R. 2021. “Reassessing the property tax.” *Available at SSRN 3800536* .
- Board, Civic Consulting Alliance. 2018. “Residential Property Assessment in Cook County: Summary of Analytical Findings.” *Report*. <https://www.ccachicago.org/wp-content/uploads/2018/02/2018-Residential-Property-Analysis-Final.pdf> .
- Bohren, J Aislinn, Peter Hull and Alex Imas. 2025. “Systemic discrimination: Theory and measurement.” *The Quarterly Journal of Economics* 140(3):1743–1799.
- Cook County Assessor’s Office. 2019. “Cook County Assessor’s Office Publicly Releases Residential Assessment Code and Models.” <https://www.cookcountyassessoril.gov/news/>

cook-county-assessors-office-publicly-releases-residential-assessment-code-and-models
Accessed 2026-02-27.

Cook County Assessor’s Office. 2020. Cook County Assessor’s COVID-19 Adjustments to Property Assessments in the South and West Suburbs. Technical report Cook County Assessor’s Office. Accessed 2026-02-27.

URL: <https://prodassets.cookcountyassessoril.gov/s3fs-public/reports/COVID19/COVIDAdjustmentsSo>

Cook County Assessor’s Office. 2021. Cook County Assessor’s COVID-19 Adjustments to 2020 Property Assessments in the North Suburbs and City of Chicago. Technical report Cook County Assessor’s Office. Accessed 2026-02-27.

URL: https://prodassets.cookcountyassessoril.gov/s3fs-public/form_documents/COVIDNorthCityTris

Cook County Assessor’s Office. 2025. “Does Location Really Matter? New Assessor’s Office Tool Shows the Sales and Data That Have the Most Impact on Your Home’s Value.” <https://www.cookcountyassessoril.gov/news/does-location-really-matter-new-assessors-office-tool-shows-sales-and-data-have-most->
Accessed 2026-02-27.

Cook County Assessor’s Office Data Department. N.d. “model-res-avm: Automated valuation model for class 200 residential properties in Cook County.” <https://github.com/ccao-data/model-res-avm>. Accessed 2026-02-27.

Cook County Open Data. 2019. “Cook County Assessor Model & Valuation Data Release.” <https://datacatalog.cookcountyil.gov/stories/s/Cook-County-Assessor-Valuation-Data-Release/p2kt-hk36/>. Accessed 2026-02-27.

Elzayn, Hadi, Evelyn Smith, Thomas Hertz, Cameron Guage, Arun Ramesh, Robin Fisher, Daniel E Ho and Jacob Goldin. 2025. “Measuring and mitigating racial disparities in tax audits.” *The Quarterly Journal of Economics* 140(1):113–163.

Friedman, Jerome H. 2001. “Greedy function approximation: a gradient boosting machine.” *Annals of statistics* pp. 1189–1232.

- Fu, Runshan, Yan Huang, Nitin Mehta, Param Vir Singh and Kannan Srinivasan. 2025. “Unequal impact of Zestimate on the housing market.” *Marketing Science* 44(6):1407–1427.
- Green, Ben. 2022. “The flaws of policies requiring human oversight of government algorithms.” *Computer Law & Security Review* 45:105681.
- Grotto, Jason, Sandhya Kambhampati and Ray Long. 2017. “The Tax Divide.”. Chicago Tribune and ProPublica Illinois investigative series.
URL: <https://features.propublica.org/the-tax-divide/cook-county-commercial-and-industrial-property-tax-assessments/>
- Holz, Justin, David Novgorodsky and Andrew Simon. 2024. “Racial Inequality in Property Tax Appeals: Evidence from Field Experiments with Homeowners and Assessors.”.
- Ihlanfeldt, Keith and Luke P Rodgers. 2021. “Explaining racial gaps in property assessment and property taxation.” *Unpublished manuscript*. Florida State University, Devoe Moore Center and Department of Economics .
- Ihlanfeldt, Keith and Luke P Rodgers. 2023. “Beyond Assessment: Racial and Gender Disparities in Property Taxation.” *Manuscript*. Florida State University, Tallahassee .
- IIAO. 2011. *Standard on mass appraisal of real property*. International Association of Assessing Officers.
- International Association of Assessing Officers. 2025. “Standard on Ratio Studies (Exposure Draft).” https://www.iaao.org/media/standards/Standard_on_Ratio_Studies_Exposure_Draft_April_2025.pdf. Accessed 2026-02-27.
- Kahrl, Andrew W. 2024. The black tax: 150 years of theft, exploitation, and dispossession in America. In *The Black Tax*. University of Chicago Press.
- Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye and Tie-Yan Liu. 2017. “Lightgbm: A highly efficient gradient boosting decision tree.” *Advances in neural information processing systems* 30.

- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig and Sendhil Mullainathan. 2018. "Human decisions and machine predictions." *The quarterly journal of economics* 133(1):237–293.
- Lambrecht, Anja and Catherine Tucker. 2019. "Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads." *Management science* 65(7):2966–2981.
- Levy, Karen, Kyla E Chasalow and Sarah Riley. 2021. "Algorithms and decision-making in the public sector." *Annual Review of Law and Social Science* 17(1):309–334.
- Novara, Marisa and Amy Khare. 2017. "Two extremes of residential segregation: Chicago's separate worlds & policy strategies for integration." *Proceedings of A Shared Future: Fostering Communities of Inclusion in an Era of Inequality* pp. 202–215.
- Quaintance, Zack. 2019. "Cook County, Ill., Opens Access to Residential Assessment Code." <https://www.govtech.com/analytics/Cook-County-Ill-Opens-Access-to-Residential-Assessment-Code.html>. Accessed 2026-02-27.
- Rosen, Sherwin. 1974. "Hedonic prices and implicit markets: product differentiation in pure competition." *Journal of political economy* 82(1):34–55.
- Wang, Daikun and Victor Jing Li. 2019. "Mass appraisal models of real estate in the 21st century: A systematic literature review." *Sustainability* 11(24):7006.
- Zhang, Shunyuan, Nitin Mehta, Param Vir Singh and Kannan Srinivasan. 2021. "Frontiers: Can an artificial intelligence algorithm mitigate racial economic inequality? An analysis in the context of Airbnb." *Marketing Science* 40(5):813–820.
- Zhang, Shuo and Peter J Kuhn. 2024. Measuring bias in job recommender systems: Auditing the algorithms. Technical report National Bureau of Economic Research.

Online Appendix

A Robustness checks to the main DiD results

Table A1: Robustness: Price Regressivity and Racial Bias — No North Triad

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.035 (0.030)		-0.029 (0.037)
Treat × Pre3 × ln(price)	0.026 (0.032)		0.019 (0.042)
Treat × Pre2 × ln(price)	0.020 (0.015)		0.018 (0.019)
Treat × Post0 × ln(price)	0.145*** (0.044)		0.168*** (0.041)
Treat × Post1 × ln(price)	0.157*** (0.031)		0.163*** (0.032)
Treat × Post2 × ln(price)	0.181*** (0.030)		0.188*** (0.033)
Treat × Post3 × ln(price)	0.195*** (0.038)		0.222*** (0.039)
Treat × Post4 × ln(price)	0.173*** (0.044)		0.194*** (0.046)
Treat × Post5 × ln(price)	0.189*** (0.034)		0.208*** (0.034)
Treat × Post6 × ln(price)	0.154*** (0.041)		0.171*** (0.040)
Treat × Post7 × ln(price)	0.173*** (0.047)		0.204*** (0.049)
Treat × Pre4 × minority		0.083* (0.044)	-0.053 (0.073)
Treat × Pre3 × minority		-0.044 (0.051)	-0.019 (0.053)
Treat × Pre2 × minority		0.025 (0.035)	0.019 (0.039)
Treat × Post0 × minority		-0.116* (0.066)	0.066 (0.044)
Treat × Post1 × minority		-0.083 (0.074)	0.058 (0.052)
Treat × Post2 × minority		-0.053 (0.065)	0.074 (0.050)
Treat × Post3 × minority		-0.082 (0.064)	0.091* (0.047)
Treat × Post4 × minority		-0.107 (0.068)	0.091 (0.059)
Treat × Post5 × minority		-0.095 (0.077)	0.080* (0.042)
Treat × Post6 × minority		-0.086 (0.075)	0.062 (0.038)
Treat × Post7 × minority		-0.047 (0.083)	0.126*** (0.041)
2020 price slope	-0.418*** (0.047)		-0.496*** (0.051)
2020 race slope		0.006 (0.030)	-0.476*** (0.083)
Observations	201,166	201,166	201,166
R-squared	0.634	0.243	0.680
Town-by-quarter FE	YES	YES	YES

Notes: This table presents robustness checks to the main DiD event study estimates (Table 1) by using only Chicago triad and the South triad (i.e., removing the North triad). Across specifications, the outcome variable is ln(Assessed/Sale Price), where Assessed is the initial assessed values. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space. All specifications include township-by-quarter fixed effects. Standard errors are clustered at the township level and in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2: Robustness: Price Regressivity and Racial Bias — No South Triad

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.024 (0.024)		0.003 (0.030)
Treat × Pre3 × ln(price)	0.056* (0.028)		0.054 (0.036)
Treat × Pre2 × ln(price)	0.001 (0.018)		-0.000 (0.023)
Treat × Post0 × ln(price)	0.168*** (0.045)		0.192*** (0.040)
Treat × Post1 × ln(price)	0.176*** (0.028)		0.179*** (0.026)
Treat × Post2 × ln(price)	0.212*** (0.026)		0.217*** (0.029)
Treat × Post3 × ln(price)	0.276*** (0.044)		0.302*** (0.044)
Treat × Pre4 × minority		0.054 (0.042)	0.004 (0.065)
Treat × Pre3 × minority		-0.126** (0.051)	-0.035 (0.048)
Treat × Pre2 × minority		-0.030 (0.058)	-0.018 (0.047)
Treat × Post0 × minority		-0.228*** (0.079)	0.037 (0.048)
Treat × Post1 × minority		-0.182** (0.075)	0.030 (0.039)
Treat × Post2 × minority		-0.209*** (0.060)	0.043 (0.038)
Treat × Post3 × minority		-0.226*** (0.070)	0.065 (0.053)
2020 price slope	-0.385*** (0.061)		-0.468*** (0.070)
2020 race slope		-0.020 (0.033)	-0.499*** (0.099)
Observations	177,972	177,972	177,972
R-squared	0.584	0.215	0.635
Town-by-quarter FE	YES	YES	YES

Notes: This table presents robustness checks to the main DiD event study estimates (Table 1) by using only Chicago triad and the North triad (i.e., removing the South triad). Across specifications, the outcome variable is ln(Assessed/Sale Price), where Assessed is the initial assessed values. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space. All specifications include township-by-quarter fixed effects. Standard errors are clustered at the township level and in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A3: Robustness: Price Regressivity and Racial Bias — Triad-Level Wild Cluster Bootstrap

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.032 (0.101)		-0.021 (0.311)
Treat × Pre3 × ln(price)	0.035 (0.322)		0.028 (0.392)
Treat × Pre2 × ln(price)	0.014 (0.194)		0.013 (0.203)
Treat × Post0 × ln(price)	0.154 (0.260)		0.177 (0.276)
Treat × Post1 × ln(price)	0.170 (0.245)		0.175* (0.213)
Treat × Post2 × ln(price)	0.197 (0.379)		0.202 (0.358)
Treat × Post3 × ln(price)	0.223 (0.778)		0.249 (0.767)
Treat × Post4 × ln(price)	0.158 (1.261)		0.178 (1.464)
Treat × Post5 × ln(price)	0.175 (1.215)		0.192 (1.344)
Treat × Post6 × ln(price)	0.139 (1.282)		0.155 (1.415)
Treat × Post7 × ln(price)	0.158 (1.462)		0.188 (1.594)
Treat × Pre4 × minority		0.075 (0.293)	-0.037 (0.589)
Treat × Pre3 × minority		-0.069 (0.889)	-0.026 (0.213)
Treat × Pre2 × minority		0.009 (0.583)	0.008 (0.418)
Treat × Post0 × minority		-0.149 (1.160)	0.057 (0.317)
Treat × Post1 × minority		-0.116 (0.989)	0.054 (0.234)
Treat × Post2 × minority		-0.102 (1.621)	0.065 (0.331)
Treat × Post3 × minority		-0.124 (1.399)	0.084 (0.249)
Treat × Post4 × minority		-0.119 (1.171)	0.062 (2.752)
Treat × Post5 × minority		-0.106 (1.024)	0.051 (2.510)
Treat × Post6 × minority		-0.098 (1.079)	0.033 (2.661)
Treat × Post7 × minority		-0.059 (1.309)	0.097 (3.096)
2020 price slope	-0.384* (0.322)		-0.452* (0.513)
2020 race slope		-0.003 (0.124)	-0.433 (0.839)
Observations	284,935	284,935	284,935
R-squared	0.604	0.237	0.647
Town-by-quarter FE	YES	YES	YES

Notes: This table presents the main DiD (Equation 1) estimates examining how the algorithmic reform affected how the assessment values vary with logged sale price (the price slope) and Census-Block-Group-level share of minority population (the race slope). Across specifications, the outcome variable is $\ln(\text{Assessed}/\text{Sale Price})$, where Assessed is the initial assessed values. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space, but they can be seen in the event-study figures. All specifications include township-by-quarter fixed effects. Wild-cluster bootstrap inference is conducted at the triad level using `boottest`, Webb weights, and 9999 replications. Because `boottest` returns p-values and confidence sets rather than a conventional variance estimator, the numbers in parentheses are implied standard errors backed out from the 95% bootstrap confidence interval. Significance stars are based on bootstrap p-values. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A4: Robustness: Price Regressivity and Racial Bias — Census-Block-level Minority

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.032 (0.024)		-0.016 (0.030)
Treat × Pre3 × ln(price)	0.035 (0.026)		0.038 (0.033)
Treat × Pre2 × ln(price)	0.014 (0.014)		0.004 (0.019)
Treat × Post0 × ln(price)	0.154*** (0.042)		0.176*** (0.042)
Treat × Post1 × ln(price)	0.170*** (0.027)		0.184*** (0.027)
Treat × Post2 × ln(price)	0.197*** (0.026)		0.200*** (0.032)
Treat × Post3 × ln(price)	0.223*** (0.038)		0.254*** (0.039)
Treat × Post4 × ln(price)	0.158*** (0.044)		0.174*** (0.045)
Treat × Post5 × ln(price)	0.175*** (0.033)		0.193*** (0.034)
Treat × Post6 × ln(price)	0.139*** (0.040)		0.158*** (0.041)
Treat × Post7 × ln(price)	0.158*** (0.046)		0.199*** (0.049)
Treat × Pre4 × minority		0.098*** (0.028)	-0.028 (0.068)
Treat × Pre3 × minority		-0.032 (0.041)	0.014 (0.044)
Treat × Pre2 × minority		0.012 (0.034)	-0.030 (0.027)
Treat × Post0 × minority		-0.118** (0.053)	0.046 (0.030)
Treat × Post1 × minority		-0.092 (0.066)	0.079** (0.034)
Treat × Post2 × minority		-0.106** (0.049)	0.051 (0.050)
Treat × Post3 × minority		-0.104* (0.055)	0.091*** (0.034)
Treat × Post4 × minority		-0.078 (0.069)	0.054 (0.052)
Treat × Post5 × minority		-0.102 (0.070)	0.035 (0.041)
Treat × Post6 × minority		-0.096 (0.064)	0.037 (0.043)
Treat × Post7 × minority		-0.006 (0.091)	0.135*** (0.043)
2020 price slope	-0.384*** (0.043)		-0.451*** (0.050)
2020 race slope		0.005 (0.024)	-0.413*** (0.071)
Observations	284,935	284,935	284,935
R-squared	0.604	0.237	0.649
Town-by-quarter FE	YES	YES	YES

Notes: This table presents robustness checks to the main DiD event study estimates (Table 1) by using the Census-Block-level share of minority population instead of the CBG-specific minority. Across specifications, the outcome variable is $\ln(\text{Assessed}/\text{Sale Price})$, where Assessed is the initial assessed values. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space. All specifications include township-by-quarter fixed effects. Standard errors are clustered at the township level and in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A5: Robustness: Price Regressivity and Racial Bias — Inferred Owner Minority from Seller Name

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.011 (0.029)		-0.006 (0.029)
Treat × Pre3 × ln(price)	0.046 (0.033)		0.049 (0.032)
Treat × Pre2 × ln(price)	0.009 (0.015)		0.011 (0.015)
Treat × Post0 × ln(price)	0.151*** (0.050)		0.151*** (0.049)
Treat × Post1 × ln(price)	0.164*** (0.034)		0.168*** (0.032)
Treat × Post2 × ln(price)	0.198*** (0.031)		0.203*** (0.028)
Treat × Post3 × ln(price)	0.217*** (0.037)		0.221*** (0.035)
Treat × Post4 × ln(price)	0.147*** (0.046)		0.149*** (0.044)
Treat × Post5 × ln(price)	0.169*** (0.040)		0.172*** (0.038)
Treat × Post6 × ln(price)	0.152*** (0.048)		0.156*** (0.047)
Treat × Post7 × ln(price)	0.156*** (0.047)		0.161*** (0.045)
Treat × Pre4 × minority		-0.048** (0.023)	-0.011 (0.021)
Treat × Pre3 × minority		0.004 (0.023)	0.013 (0.021)
Treat × Pre2 × minority		0.016 (0.034)	0.032* (0.017)
Treat × Post0 × minority		-0.017 (0.023)	-0.001 (0.011)
Treat × Post1 × minority		0.042 (0.026)	0.065*** (0.023)
Treat × Post2 × minority		0.011 (0.030)	0.039 (0.024)
Treat × Post3 × minority		-0.020 (0.031)	0.009 (0.021)
Treat × Post4 × minority		0.017 (0.029)	0.031 (0.020)
Treat × Post5 × minority		0.002 (0.024)	0.021 (0.016)
Treat × Post6 × minority		-0.002 (0.025)	0.019 (0.025)
Treat × Post7 × minority		0.000 (0.033)	0.039 (0.030)
2020 price slope	-0.348*** (0.043)		-0.349*** (0.043)
2020 race slope		-0.016 (0.013)	-0.029** (0.013)
Observations	247,618	247,618	247,618
R-squared	0.516	0.182	0.520
Town-by-quarter FE	YES	YES	YES

Notes: This table presents robustness checks to the main DiD event study estimates (Table 1) by using the inferred owner minority probability using their names. Across specifications, the outcome variable is $\ln(\text{Assessed}/\text{Sale Price})$, where Assessed is the initial assessed values. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space. All specifications include township-by-quarter fixed effects. Standard errors are clustered at the township level and in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A6: Robustness: Price Regressivity and Racial Bias — Neighborhood FE and Neighborhood-specific Slopes

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.044* (0.023)		-0.043* (0.023)
Treat × Pre3 × ln(price)	0.015 (0.022)		0.015 (0.022)
Treat × Pre2 × ln(price)	0.001 (0.018)		0.002 (0.018)
Treat × Post0 × ln(price)	0.101*** (0.027)		0.102*** (0.027)
Treat × Post1 × ln(price)	0.097*** (0.024)		0.096*** (0.023)
Treat × Post2 × ln(price)	0.121*** (0.024)		0.121*** (0.024)
Treat × Post3 × ln(price)	0.199*** (0.023)		0.201*** (0.023)
Treat × Post4 × ln(price)	0.149*** (0.028)		0.153*** (0.028)
Treat × Post5 × ln(price)	0.143*** (0.027)		0.148*** (0.027)
Treat × Post6 × ln(price)	0.089*** (0.030)		0.091*** (0.030)
Treat × Post7 × ln(price)	0.124*** (0.026)		0.126*** (0.026)
Treat × Pre4 × minority		0.134 (0.088)	0.051 (0.060)
Treat × Pre3 × minority		0.036 (0.089)	-0.021 (0.056)
Treat × Pre2 × minority		0.112 (0.075)	0.071 (0.048)
Treat × Post0 × minority		0.116 (0.090)	0.077 (0.066)
Treat × Post1 × minority		0.042 (0.089)	-0.003 (0.063)
Treat × Post2 × minority		0.072 (0.086)	0.062 (0.059)
Treat × Post3 × minority		0.095 (0.093)	0.105 (0.069)
Treat × Post4 × minority		-0.125 (0.106)	0.002 (0.068)
Treat × Post5 × minority		-0.049 (0.108)	0.017 (0.066)
Treat × Post6 × minority		0.033 (0.116)	0.077 (0.066)
Treat × Post7 × minority		0.015 (0.121)	0.054 (0.077)
2020 price slope	-0.655***(0.016)		-0.656***(0.016)
2020 race slope		0.033*(0.019)	-0.081***(0.014)
Observations	282,584	282,584	282,584
R-squared	0.797	0.367	0.800
Nbhd-by-quarter FE	YES	YES	YES

Notes: This table presents robustness checks to the main DiD event study estimates (Table 1) by allowing for neighborhood fixed effects and neighborhood-specific price and minority slopes. Across specifications, the outcome variable is $\ln(\text{Assessed}/\text{Sale Price})$, where Assessed is the initial assessed values. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space. All specifications include township-by-quarter fixed effects. Standard errors are clustered at the township level and in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A7: Robustness: Price Regressivity and Racial Bias — Census-Tract level Gentrification

	(1) Price slope	(2) Race slope	(3) Price slope + race slope
Treat × Pre4 × ln(price)	-0.024 (0.023)		-0.015 (0.032)
Treat × Pre3 × ln(price)	0.044 (0.028)		0.040 (0.037)
Treat × Pre2 × ln(price)	0.022* (0.012)		0.020 (0.017)
Treat × Post0 × ln(price)	0.154*** (0.038)		0.181*** (0.033)
Treat × Post1 × ln(price)	0.168*** (0.028)		0.175*** (0.026)
Treat × Post2 × ln(price)	0.182*** (0.027)		0.191*** (0.028)
Treat × Post3 × ln(price)	0.225*** (0.043)		0.250*** (0.039)
Treat × Post4 × ln(price)	0.160*** (0.044)		0.181*** (0.043)
Treat × Post5 × ln(price)	0.171*** (0.033)		0.192*** (0.032)
Treat × Post6 × ln(price)	0.135*** (0.038)		0.154*** (0.036)
Treat × Post7 × ln(price)	0.142*** (0.043)		0.178*** (0.045)
Treat × Pre4 × minority		0.047 (0.039)	-0.048 (0.070)
Treat × Pre3 × minority		-0.070* (0.040)	-0.032 (0.051)
Treat × Pre2 × minority		-0.004 (0.034)	0.006 (0.038)
Treat × Post0 × minority		-0.106* (0.055)	0.082** (0.035)
Treat × Post1 × minority		-0.075 (0.060)	0.058 (0.041)
Treat × Post2 × minority		-0.069 (0.056)	0.070* (0.038)
Treat × Post3 × minority		-0.076 (0.058)	0.100** (0.038)
Treat × Post4 × minority		-0.057 (0.068)	0.088 (0.060)
Treat × Post5 × minority		-0.053 (0.066)	0.079* (0.042)
Treat × Post6 × minority		-0.045 (0.069)	0.061 (0.041)
Treat × Post7 × minority		0.023 (0.072)	0.139*** (0.042)
2020 price slope	-0.391*** (0.043)		-0.460*** (0.048)
2020 race slope		-0.000 (0.026)	-0.440*** (0.071)
Observations	273,551	273,551	273,551
R-squared	0.614	0.243	0.655
Town-by-quarter FE	YES	YES	YES

Notes: This table presents robustness checks to the main DiD event study estimates (Table 1) by allowing for baseline, slopes, and effects to vary across Census-tract level gentrification (gentrification-related estimates not reported here since they are not of the primary interest). Across specifications, the outcome variable is $\ln(\text{Assessed/Sale Price})$, where Assessed is the initial assessed values. The reference period is $\tau = -1$ (2020 Q4). We do not report here the dynamic treatment effects for the periods before -4 to save space. All specifications include township-by-quarter fixed effects. Standard errors are clustered at the township level and in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.