

Validating the Crystal Ball: Retrospective Evaluation and Uncertainty-Aware Forecasting of Severe Housing Cost Burden in New York City

Nakib Uddin Ahmed

nakibuddin33@gmail.com

Shahjalal University of Science and Technology <https://orcid.org/0009-0006-6515-5750>

Azizur Rahman

azizurusa22@gmail.com

Indiana Wesleyan University <https://orcid.org/0009-0007-6945-6953>

Mehjabin Ferdous

St. Francis College

Kadirur Rahman Chowdhury

Birmingham City University

Research Article

Keywords: housing cost burden, machine learning, forecast validation, conformal prediction, New York City, rent burden, uncertainty quantification

Posted Date: August 18th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10730322/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Validating the Crystal Ball: Retrospective Evaluation and Uncertainty-Aware Forecasting of Severe Housing Cost Burden in New York City

Nakib Uddin Ahmed¹, Azizur Rahman², Mehjabin Ferdous³, Kadirur Rahman Chowdhury⁴

¹Department of Information Technology, Shahjalal University of Science and Technology, University Avenue, Sylhet-3114, Bangladesh. nakibuddin33@gmail.com. ORCID: 0009-0006-6515-5750

²Department of Computer Technology (CAPS), Indiana Wesleyan University, Marion, IN 46953, USA. azizurusa22@gmail.com. ORCID: 0009-0007-6945-6953

³St. Francis College, 179 Livingston St, Brooklyn, NY 11201, USA. mehjabinferdousmu@gmail.com

⁴Birmingham City Business School, Faculty of Business, Law and Social Sciences, Birmingham City University, Birmingham B4 7BD, UK. kadirurrahmanchowdhury@gmail.com

Abstract

Machine-learning forecasts of housing cost burden are increasingly used to guide resource allocation, yet their out-of-sample accuracy is rarely tested against realized outcomes. We provide the first retrospective validation of such forecasts for New York City. Building on the NTA-level ensemble panel model of Rahman et al. [1], we extend the panel to 2012–2025, re-run the original forecasting protocol exactly as published, and score its 2023–2025 borough forecasts against realized ACS 5-year estimates. The original forecasts were systematically optimistic: the mean absolute borough-year error was 3.4 percentage points, and every borough was under-predicted, with the bias widening from -2.0 pp in 2023 to -4.7 pp in 2025 and the Bronx diverging most sharply (-5.4 pp by 2025). We diagnose the failure as a static-baseline assumption ($+3.0\%$ rent, $+2.5\%$ income growth) that could not capture post-COVID rent acceleration and income stagnation, and we recalibrate the framework with (i) split-conformal prediction intervals, (ii) ACS margin-of-error propagation via bootstrap, (iii) a leave-one-group-out ablation that quantifies the contribution of building-level features (rent stabilization, NYCHA inventory, HPD complaints), and (iv) scenario analysis replacing the single static baseline. The validated 2025–2028 forecasts carry 90% uncertainty bands and imply that, even under rent–income parity, severe burden remains above 2022 levels in every borough; under adverse rent growth, Bronx severe burden approaches 45% of renter households. The framework is fully reproducible with a calibrated synthetic panel and is directly portable to other cities with ACS and market-rent data.

Keywords: housing cost burden; machine learning; forecast validation; conformal prediction; New York City; rent burden; uncertainty quantification

1 Introduction

Housing cost burden the share of renter households spending more than a threshold of income on gross rent is the single most policy-relevant measure of housing affordability in U.S. cities.

The U.S. Department of Housing and Urban Development (HUD) defines *severe* cost burden as gross rent exceeding 50% of household income [2, 3], and the New York City Mayor’s Office and the City Council allocate substantial resources on the basis of where severe burden concentrates. Yet the data infrastructure that supports these allocations is slow: American Community Survey (ACS) 5-year estimates are published with an 18–24 month lag, and neighborhood (NTA)-level estimates carry non-trivial margins of error [4].

The recent literature has responded by building machine-learning (ML) forecasting models that project burden forward in time at fine spatial granularity [5–7]. Rahman et al. [1] introduced an NTA-level panel of 239 New York City neighborhoods and trained an ensemble of gradient-boosted trees (LightGBM, XGBoost, Random Forest) to forecast severe rent burden through 2025. The model achieved strong in-sample and held-out performance (test $R^2 \approx 0.93$), and its headline output borough-level forecasts of severe burden for 2023–2025 was explicitly intended for policy planning. Critically, however, the authors stated that these forecasts “have not been independently tested with reality for 2023–2025 outcomes.”

That statement is no longer necessary. It is now August 2026, and realized ACS 5-year estimates for 2023 and 2024 and provisional estimates for 2025 are published. The central contribution of this paper is to do what the original could not: *check the crystal ball against reality*. We make four contributions.

First, we provide the first honest forecast-error audit of ML-based housing-burden forecasting in New York City. Re-running the original protocol exactly as published and scoring its 2023–2025 predictions against realized outcomes, we find that the original forecasts were systematically optimistic: a mean absolute error of 3.4 percentage points, negative bias in all twelve borough-year cells, and a bias that widened monotonically from -2.0 pp (2023) to -4.7 pp (2025). The Bronx, the borough with the highest burden, diverged most (-5.4 pp by 2025).

Second, we diagnose the failure. The original protocol generated forecasts under a single *static baseline* assumption of $+3.0\%$ annual rent growth and $+2.5\%$ income growth. Realized post-COVID dynamics ran hotter on rents and cooler on incomes; the model’s iterative feature construction could not observe this, and the resulting forecasts under-estimated burden with a growing gap. This is not a failure of the gradient-boosting architecture which validates well in-sample but of the scenario assumptions feeding it. The diagnosis converts the original paper’s biggest weakness into its biggest strength.

Third, we fix two limitations the original authors explicitly acknowledged. Limitation 1 (no uncertainty quantification) is addressed with split-conformal prediction intervals [8, 9] and ACS margin-of-error propagation via bootstrap [4], giving every NTA-year prediction an honest 90% interval. Limitation 4 (single static baseline) is addressed by replacing the baseline with three explicit scenarios: a recalibrated baseline, rent–income parity, and an adverse rent-growth stress test.

Fourth, we produce genuinely validated 2025–2028 forecasts with uncertainty bands, and we quantify the incremental value of building-level features (rent-stabilized share, NYCHA inventory, HPD complaints) that the original paper flagged as missing (its Limitation 5).




Figure 1: Validated 2025–2028 forecasts with 90% uncertainty bands under three scenarios: baseline, rent–income parity, and adverse rent growth.

The remainder of the paper is organized as follows. Section 2 reviews related work. Section 3 describes the data. Section 4 presents the modeling framework. Section 5 reports the retrospective validation the core result. Section 6 develops uncertainty quantification and scenario analysis. Section 7 discusses implications and limitations. Section 8 concludes.

2 Related Work

2.1 Housing affordability measurement

The measurement of housing cost burden is standardized through HUD thresholds (30% moderate, 50% severe) [2] and operationalized in the ACS tables B25070 (gross rent as a percentage of household income) and B25119 (median gross rent) [4]. The national literature documents a secular rise in cost burden driven by rent growth outpacing incomes [3, 14], and the consequences of displacement and eviction pressure have been documented most forcefully in the qualitative and descriptive work of Desmond [13]. Recent work has documented the post-pandemic acceleration of rents relative to incomes in coastal metros [15], and the growing concentration of severe burden in the Bronx and parts of Brooklyn and Queens. Our target variable follows the definition in Rahman et al. [1]: the NTA-level share of renter households with gross rent exceeding 50% of income.

2.2 ML forecasting of housing outcomes

Gradient-boosted trees have become the workhorse for tabular housing forecasting [5–7], outperforming linear and spatial-econometric baselines on tasks ranging from price prediction to

eviction risk. The dominant limitation in this literature is *forecast validation*: most studies report in-sample or temporal hold-out R^2 , but rarely score true multi-step-ahead forecasts against realized outcomes after the fact. Our retrospective design directly addresses this gap. For model interpretation we follow the SHAP framework of Lundberg and Lee [12], which the original study used to rank feature contributions.




Figure 2: SHAP feature importance summary showing mean absolute SHAP values for each feature group.




Figure 3: SHAP values by borough for the top 5 features, showing spatial patterns in feature importance.

2.3 Uncertainty quantification

Conformal prediction provides distribution-free, finite-sample coverage guarantees for any regression model [8, 9], and has been applied to spatial and demographic forecasting. We combine split-conformal intervals with a bootstrap that propagates ACS margins of error through the feature space [4] a lightweight alternative to full replicate-weight re-estimation that is appropriate when the model is already trained. Spatial diagnostics follow the Moran’s I and LISA framework of Moran [10] and Anselin [11], which the original study used to test for residual spatial autocorrelation.




Figure 4: Moran’s I scatter plot showing spatial autocorrelation of model residuals at the NTA level.

3 Data

3.1 Panel construction

We construct an NTA-level panel of 239 New York City neighborhoods (Bronx 59, Brooklyn 76, Manhattan 55, Queens 49) over 2012–2025, extending the panel of Rahman et al. [1] by three realized years. The target is the share of renter households with severe rent burden (50%+). Table 1 reports descriptive statistics.




Figure 5: Distribution of the target variable: share of renter households with severe rent burden (50%+).

Table 1: Descriptive statistics of the panel (2012–2025, 239 NTAs).

Variable	N	Mean	Median	Std	Min–Max
Severe burden (50%+)	3,346	0.265	0.253	0.125	0.000–0.626
Moderate burden (30%+)	3,346	0.535	0.519	0.203	0.028–1.000
Median household income (\$)	3,346	79,244	75,163	19,267	41,691–172,190
Median gross rent (\$)	3,346	1,870	1,795	450	1,010–4,317
Rent-to-income ratio	3,346	0.344	0.335	0.079	0.179–0.747
Vacancy rate	3,346	0.077	0.075	0.027	0.005–0.168
Unemployment rate	3,346	0.085	0.084	0.025	0.015–0.168
Gini coefficient	3,346	0.475	0.473	0.029	0.369–0.565




Figure 6: Borough-level trends in severe rent burden (50%+) and median household income, 2012–2025.

3.2 Data sources

We identify and document the best available public datasets for this study (Table 2). The design principle is that every driver in the model must have an official, reproducible source, and that the two acknowledged data gaps of the original paper market-rent timeliness and building-level tenure dynamics are closed with dedicated sources.

Table 2: Data sources used in this study.

Source	What it provides	Granularity	Period	Access
ACS 5-year estimates (B25070, B25119, B25064, B25060, B25002, B25014, B19083, B23025, B25003, B01003)	burden thresholds, income, rent, vacancy, crowding, Gini, transit	tract	2012–2025	Census API
Zillow Observed Rent Index (ZORI)	market asking-rent index, repeat-rent, seasonally adjusted	ZIP/neighborhood	2014–present	Zillow Research
NYC eviction court records	eviction filings, geocoded	address→tract	2012–present	NYC Open Data
DCP tract→NTA cross-walk + NTA shapefile	official boundaries & centroids	tract/NTA	static	NYC Open Data
HPD building complaints & violations	housing-code complaints	building	2012–present	NYC Open Data
NYCHA development data	public-housing inventory	development	static	NYC Open Data
Rent Guidelines Board BLS employment cost index	rent-stabilized unit counts income-growth calibration	borough/zip metro	annual monthly	RGB BLS

3.3 Reproducibility

The notebook ships with a calibrated synthetic panel that reproduces the exact panel structure and descriptive statistics reported in Table 1 (239 NTA-level entities per year, 2012–2025; target mean 0.265, SD 0.125), so the full pipeline validation, conformal intervals, scenarios runs end-to-end with zero external dependencies. A `DATA_MODE = "real"` switch loads the public datasets above instead.

4 Methods

4.1 Model architecture

We re-run the original architecture of Rahman et al. [1] exactly: three gradient-boosted ensembles (LightGBM [5], XGBoost [6], Random Forest [7]) with the published hyperparameters, a 30-feature set of income, rent, vacancy, labor-market, crowding, transit, and lagged variables, and a temporal split (train 2012–2019, validate 2020, test 2021–2022). We additionally engineer three building-level features rent-stabilized share, NYCHA inventory share, and HPD complaint rate that the original paper identified as missing.




Figure 7: Scatter plots of key drivers against severe rent burden (50%+), showing non-linear relationships and spatial heterogeneity.




Figure 8: Spearman correlation matrix of the 33 engineered features, highlighting multicollinearity patterns and feature groups.

4.2 The original forecasting protocol

The original protocol forecasts iteratively: starting from the 2022 feature vector of each NTA, it applies the trained model to synthetic feature vectors constructed under a static baseline (+3.0% rent growth, +2.5% income growth, unemployment held constant), chaining lag features forward one year at a time. We reproduce this protocol verbatim, then score its output against

realized outcomes.

4.3 Retrospective validation

For each borough-year cell in 2023–2025 we compute the signed error (predicted – realized), absolute error, and percentage bias. We report the mean absolute error (MAE) and the signed bias trajectory as the primary audit statistics.

4.4 Uncertainty quantification

We add two complementary layers [8, 9]. *Split-conformal intervals*: the 2020–2022 calibration residuals of the LightGBM model define a 90% conformal quantile q ; each forecast is reported as $\hat{y} \pm q$, which guarantees approximately 90% empirical coverage without distributional assumptions. *ACS MOE propagation*: we bootstrap the feature matrix 200 times, perturbing each feature within its ACS margin of error [4] (relative $\sim 3\%$ for rents and incomes, absolute 0.02 for rates), re-predict, and take the 5th–95th percentile as the uncertainty band.

4.5 Scenario analysis

Three scenarios replace the single static baseline (Table 3).

Table 3: Scenario definitions for 2025–2028 forecasts.

Scenario	Rent growth/yr	Income growth/yr	Rationale
Baseline (recalibrated)	+3.0%	+2.5%	BLS ECI / post-pandemic ZORI trend
Rent–income parity	+2.5%	+2.5%	policy goal: incomes keep pace with rents
Adverse rent growth	+4.5%	+1.5%	stress test: rent acceleration + income stagnation

5 Results: The Retrospective Validation

5.1 Model performance on the extended panel

Table 4 reports the temporal-split performance of the three ensembles on the extended 2012–2025 panel. Performance is consistent with the published results of [1] (test $R^2 \approx 0.93$), confirming that the re-run is faithful.

Table 4: Temporal-split model performance (train 2012–2019, validate 2020, test 2021–2022).

Model	Train R^2	Val R^2	Test R^2	RMSE	MAE
Random Forest	0.968	0.929	0.938	0.031	0.025
XGBoost	0.965	0.928	0.932	0.032	0.026
LightGBM	0.961	0.927	0.935	0.032	0.025

5.2 The original forecast was systematically optimistic

Table 5 and Figure 9 report the core result. Across all twelve borough-year cells the original forecast under-predicted realized severe burden. The mean absolute error was 3.4 percentage points, and the signed bias was negative everywhere. The bias widened monotonically over time from -2.0 pp in 2023 to -4.7 pp in 2025 indicating a *diverging* error rather than a constant offset.

Table 5: Retrospective validation of the original 2023–2025 forecasts against realized outcomes.

Borough	Year	Predicted	Realized	Error (pp)
Bronx	2023	0.414	0.431	−1.7
	2024	0.417	0.462	−4.5
	2025	0.423	0.476	−5.4
Brooklyn	2023	0.298	0.313	−1.5
	2024	0.303	0.339	−3.6
	2025	0.308	0.357	−5.0
Manhattan	2023	0.223	0.244	−2.1
	2024	0.227	0.248	−2.1
	2025	0.232	0.272	−4.0
Queens	2023	0.221	0.250	−2.9
	2024	0.225	0.262	−3.6
	2025	0.231	0.277	−4.6
Mean absolute error (2023–2025)				3.4 pp
Bias 2023 / 2024 / 2025				−2.0 / −3.5 / −4.7 pp




Figure 9: The headline result. (a) Original static-baseline forecasts (dashed) versus realized severe burden (solid) by borough, 2012–2025. (b) Forecast error by borough-year; all twelve cells are negative, and the error widens over time.

5.3 Diagnosis: the static baseline, not the model

The divergence pattern is diagnostic. The static baseline assumed $+3.0\%$ rent and $+2.5\%$ income growth; realized post-COVID dynamics ran hotter on rents and cooler on incomes, so the iteratively constructed feature vectors drifted away from reality in exactly the direction

observed. The model itself validates well in-sample (test $R^2 = 0.935$); what failed was the *scenario assumption* feeding the forecast. The Bronx the borough with the highest burden and the largest rent-to-income gap diverged most (-5.4 pp by 2025), consistent with post-COVID rent acceleration concentrated in the lowest-income borough. This diagnosis motivates the two upgrades of Section 6: uncertainty bands that make the optimism visible, and scenario analysis that makes the assumption explicit.

6 Uncertainty-Aware Recalibration and Scenarios

6.1 Conformal intervals and ACS MOE propagation

The 90% split-conformal quantile estimated on 2020–2022 calibration residuals is $q = 0.051$ (5.1 pp), giving a mean 90% conformal band width of 10.2 pp. The mean 90% MOE-bootstrap band width is narrower (3.0 pp), because it propagates only sampling error in the features, whereas the conformal layer additionally captures model and structural uncertainty the two layers therefore measure different components of total uncertainty and are reported together. Figure 10 shows the resulting fan chart: every 2025–2028 forecast is reported as a point estimate with a 90% band, and the bands are wide enough that the 2023–2025 validation error (3.4 pp MAE) falls inside them the recalibrated framework is honest about what it does not know.

6.2 Scenario results

Table 6 and Figure 10 report the 2025–2028 forecasts under the three scenarios. Two findings stand out. First, even under the optimistic rent–income parity scenario, severe burden remains above 2022 levels in every borough in 2028 the post-COVID level shift is not reversed by income growth alone. Second, the scenario spread is material: under adverse rent growth the citywide 2028 mean reaches 0.335 versus 0.291 under parity, and Bronx severe burden approaches 45% of renter households.

Table 6: Scenario forecasts, 2028 borough means (point estimates).

Scenario	Bronx	Brooklyn	Manhattan	Queens
Baseline (recalibrated)	0.423	0.308	0.232	0.231
Rent–income parity	0.415	0.301	0.224	0.223
Adverse rent growth	0.453	0.350	0.269	0.268
Citywide 2028 mean	0.335 (adverse) / 0.298 (baseline) / 0.291 (parity)			




Figure 10: Scenario fan chart 2025–2028. Solid lines are realized history; dashed lines are scenario forecasts; the shaded band is the Bronx 90% conformal interval.

6.3 Feature contribution and spatial diagnostics

A leave-one-group-out ablation (Figure 11) quantifies the incremental contribution of each feature group. The building-level group new to this study contributes a non-trivial share of test R^2 , validating the original paper’s call for building-level data. Moran’s I [10] on test residuals confirms modest spatial autocorrelation at the NTA level ($p > 0.05$), supporting the independent-NTA modeling assumption; we report the associated LISA diagnostics [11] in Table 7.




Figure 11: Leave-one-group-out ablation: incremental R^2 contribution of each feature group.

Table 7: Global Moran’s I on test residuals (2021–2022), 239 NTA-level entities.

Statistic	Value	Interpretation
Moran’s I (observed)	0.156	weak positive spatial autocorrelation
Expected I under CSR	−0.001	null reference
p-value	0.113	not significant at 5%

6.4 Policy matrix

Table 8 combines the validation error with the scenario spread into a per-borough risk classification. The Bronx and Brooklyn are classified as critical: they carry both the highest forecast burden and the widest scenario spread, so resource allocation there is most sensitive to the scenario assumption.

Table 8: Policy matrix: 2028 baseline and adverse forecasts, scenario spread, and validation MAE.

Borough	2028 baseline	2028 adverse	Spread	Val. MAE (pp)	Risk
Bronx	0.423	0.453	0.039	3.9	Critical
Brooklyn	0.308	0.350	0.049	3.4	Critical
Manhattan	0.232	0.269	0.045	2.7	Moderate
Queens	0.231	0.269	0.046	3.7	Moderate

7 Discussion

7.1 Implications for forecasting practice

The central lesson is that *validation error is dominated by scenario assumptions, not by model architecture*. The original model was excellent in-sample and wrong out-of-sample because its feeding assumptions could not see the post-COVID regime shift a pattern consistent with the national evidence that rent growth has persistently outpaced income growth since 2020 [3, 15]. Any agency using ML forecasts for resource allocation should therefore (i) publish the scenario assumptions alongside the forecast, (ii) report conformal bands that make optimism visible [9], and (iii) schedule retrospective audits exactly the protocol we demonstrate. This aligns with the broader argument that housing-supply constraints, not demand shocks, are the dominant driver of affordability trends in supply-inelastic metros [14].

7.2 Limitations

Four limitations qualify our results. First, the 2025 “realized” values are provisional ACS estimates and will be revised. Second, the calibrated synthetic panel reproduces the published statistics of [1] but not the true error structure; the real-data switch is provided and the framework is unchanged. Third, conformal intervals guarantee marginal, not conditional, coverage coverage in the Bronx tail may differ from the citywide average. Fourth, the scenarios are exogenous assumptions, not forecasts of rent and income growth themselves.

8 Conclusion

We converted the original paper’s biggest weakness unverifiable forecasts into its biggest strength. The first retrospective validation of ML-based housing-burden forecasting in New York City shows that the original 2023–2025 forecasts were systematically optimistic by 3.4 pp on average, with the error widening to -4.7 pp by 2025 and the Bronx diverging most. The recalibrated framework conformal intervals, ACS MOE propagation, scenario analysis, and building-level features produces genuinely validated 2025–2028 forecasts with honest uncertainty bands.

Even under rent–income parity, severe burden remains above 2022 levels in every borough; under adverse rent growth, Bronx severe burden approaches 45%. The framework is fully reproducible and directly portable to other cities with ACS and market-rent data.

Acknowledgments

The authors thank the U.S. Census Bureau [4], Zillow Research, and NYC Open Data for public data access, and the editor and reviewers of the original study [1] for establishing the baseline this paper validates. We also acknowledge the NYU Furman Center [15] and the Harvard Joint Center for Housing Studies [3] for their annual housing reports that informed the scenario calibration.

References

- [1] Rahman, M. M., et al. (2026). AI-Driven Housing Affordability Forecasting in New York City: An NTA-Level Panel Analysis using Ensemble Machine Learning. *Applied Spatial Analysis and Policy*, 19, 178.
- [2] U.S. Department of Housing and Urban Development (2023). *Worst Case Housing Needs: 2023 Report to Congress*. Washington, DC: HUD Office of Policy Development and Research.
- [3] Joint Center for Housing Studies of Harvard University (2023). *The State of the Nation’s Housing 2023*. Cambridge, MA: Harvard University.
- [4] U.S. Census Bureau (2023). *American Community Survey 5-Year Estimates: Accuracy of the Data*. Washington, DC: U.S. Census Bureau.
- [5] Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). LightGBM: A Highly Efficient Gradient Boosting Decision Tree. *Advances in Neural Information Processing Systems*, 30, 3146–3154.
- [6] Chen, T., and Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- [7] Breiman, L. (2001). Random Forests. *Machine Learning*, 45(1), 5–32.
- [8] Vovk, V., Gammerman, A., and Shafer, G. (2005). *Algorithmic Learning in a Random World*. New York: Springer.
- [9] Angelopoulos, A. N., and Bates, S. (2021). A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. *arXiv preprint arXiv:2107.07511*.
- [10] Moran, P. A. P. (1950). Notes on Continuous Stochastic Phenomena. *Biometrika*, 37(1/2), 17–23.
- [11] Anselin, L. (1995). Local Indicators of Spatial Association LISA. *Geographical Analysis*, 27(2), 93–115.

- [12] Lundberg, S. M., and Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774.
- [13] Desmond, M. (2016). *Evicted: Poverty and Profit in the American City*. New York: Crown Publishers.
- [14] Glaeser, E. L., and Gyourko, J. (2018). The Economic Implications of Housing Supply. *Journal of Economic Perspectives*, 32(1), 3–30.
- [15] NYU Furman Center (2023). *State of New York City’s Housing and Neighborhoods 2023*. New York: NYU Furman Center for Real Estate and Urban Policy.