

When a Deed Is Not a Market Sale: Foreclosure Transfers, Statutory Consideration, and Repeat-Sales House-Price Measurement

Pablo Loschi, independent researcher, Berlin, Germany. ORCID 0009-0004-9455-4713. Corresponding author: loschi.pablo@gmail.com

Working paper version 1.0.10, 14 October 2026

Abstract New York City’s public sales files can retain foreclosure-related lender acquisitions at legal consideration amounts. A party-based rule flags 16,655 potentially affected deeds among 1.47 million records from 2003 to 2025. Conventional screens retain most of them, and 94-95% of screened lender acquisitions recorded since 2016 match the public sales file at the same amount. Removing flagged endpoints changes a house repeat-sales index by 3.69 log points in 2010Q3 (simultaneous 95% band: 1.59 to 5.78), and by 2.63 under a narrower definition. The estimated peak-to-trough contrast is -0.64 points, with an interval from -3.03 to 1.25 . A hedonic comparison on the common attribute sample is consistent with pairing contributing to the later repeat-sales shortfall. Whole-path inference, alternative endpoint rules and assignments of unresolved reviewed cases qualify the result; margins over deletion placebos depend on their design. Of 400 AI-reviewed deeds, 84 remain unresolved. Agreement among resolved cases is not independent human validation. The contribution is a reproducible public-record screening application and a measured index-path difference, conditional on classifications and estimator choices.

1. Introduction

Studies of house prices drop transfers that are not arm’s-length sales, such as gifts, transfers between relatives and reposessions by lenders. Commercial data vendors flag transaction types for this purpose; Kermani and Wong (2026), for example, keep only the arm’s-length transactions in vendor deed data, which excludes deeds in lieu of foreclosure. Careful studies also make the separation by hand: Campbell et al. (2011) treat the lender’s resale after an unsuccessful foreclosure auction as the forced sale. Researchers who work from free public files have only what the record shows. Repeat-sales indices in the tradition of Bailey et al. (1963) and Case and Shiller (1987), hedonic models, studies of foreclosure discounts, and local price series all begin from a registry field labelled “consideration” or “sale price”. That field is produced by an administrative process tied to transfer-tax filing, not by an effort to measure market prices. Where tax law defines consideration differently from a negotiated price, the registry amount may inherit the difference.

This paper documents one such case. Under New York Tax Law §1401(d) and New York City Administrative Code §11-2101(9), consideration can include “the cancellation or discharge of an indebtedness or obligation”. When a lender takes title to a mortgaged property, whether as plaintiff at its own foreclosure sale, through a deed in lieu of foreclosure or by a transfer between servicer and investor, the amount

recorded in New York City’s Automated City Register Information System (ACRIS) can reflect a bid or debt instead of a price agreed by a buyer and seller. The lender’s later resale is retained as a potential market transaction, separately from its acquisition. A repeat-sales estimator that pairs the two amounts can treat a non-price as a price.

These transfers differ from the ones that screens for non-arm’s-length sales are designed to catch. A sale between relatives is usually recorded at a reduced or nominal price. A lender’s consideration is a legal quantity that typically falls in the range of market prices: it sits close to the Department of Finance market value, and 62% of these deeds survive a \$10,000 minimum and full-interest, single-parcel and residential-class filters. In the city’s free Department of Finance sales file, a standard source for New York City price series (e.g., Schwartz et al. 2003; Ellen et al. 2007), 94–95% of those recorded since 2016 appear at the recorded amount with no transfer-type flag. A researcher using that file cannot remove them without classifying each deed by its parties, which is what this paper’s classifier does.

The problem surfaced in a pilot for a companion study of financing-associated price gaps (Loschi 2026b). Among that study’s 8,208 one-to-three-family repeat pairs from 2016–2025, 232 had a lender as grantee at the first sale. These pairs fell in nominal terms 47.8% of the time in a rising market, against 8.5% of other pairs, and their first-sale amounts were multiples of \$1,000 22% of the time, against 90%. The pilot is reproducible from the seed/ folder of this paper’s replication package. This paper takes that observation to every deed recorded in four boroughs from 2003 to 2025, every residential property type conveyed by deed, and the full 2008 cycle.

That companion study (Loschi 2026b) finds that a repeat-sales cash–mortgage price contrast also depends on who transacts, including possible related-party transfers. Like this paper, it reads a registry field only after classifying deeds by their parties, but it measures a different field, and neither paper’s document review validates the other’s labels. A second companion paper (Loschi 2026a) measures statutory transfer and mortgage recording tax liability and how transactions shifted at New York’s 2019 transfer-tax thresholds. This paper draws on transfer-tax law only to interpret what a recorded consideration amount means, not to estimate any response to taxation.

The contribution is a public-record measurement application: a reproducible party-based screen and a quantified index-path comparison. Careful foreclosure studies already distinguish lender acquisition from resale (Campbell et al. 2011). Recorded-price adjustments can also arise from financing arrangements (Ben-David 2011), while changing property quality affects repeat-sales indices (Nowak and Smith 2020). The present screen addresses whether a consideration field can be interpreted as a negotiated price. Removing flagged deeds changes the transaction population; it does not correct the other mechanisms or recover a verified constant-quality counterfactual.

First, we measure incidence. We identify 16,655 deed-family documents in events for which the law can set the amount, 1.1% of deeds but 3.7–4.2% of priced deeds in 2008–09. Seven in eight are 1–3 family houses, and most are in Queens, Brooklyn and the Bronx. Their share of priced deeds rose almost ninefold from 2005–06 to 2009 and collapsed in 2011–12, coinciding with new court procedures for residential foreclosures. Since then it has ranged from 0.1% (2021, during the pandemic foreclosure moratorium) to 1.7%.

Second, we show that these amounts do not behave like negotiated prices, using evidence internal and external to the registry. Negotiated prices tend to be round; debt balances and judgment-based bids need not be. Legal-rule amounts carry cents 70% of the time and are multiples of \$1,000 18% of the time; ordinary sales carry cents 6% of the time and are multiples of \$1,000 81% of the time. They sit close to the

Department of Finance (DOF) market value and to the value implied by the parcel’s previous sale, but 24–34 log points above the value implied by the parcel’s next sale, which is almost always the lender’s resale; part of that gap is a forced-sale discount the registry cannot separate. Because their distribution overlaps ordinary transaction prices, threshold filters retain them.

Third, we show that conventional screens do not eliminate them. Besides the 62% that survive standard filters, 55% end up in a screened repeat-sales pair, twice the rate of ordinary deeds, which fits the quick resale of REO inventory.

Fourth, we quantify the consequence for a standard index. Some lender acquisitions L have both an earlier market purchase A and a later resale B . When the amount at L exceeds a comparable market price, the pair $A \rightarrow L$ can overstate appreciation into the take-title date, while $L \rightarrow B$ can overstate depreciation into the resale date. In a quarterly index of 1–3 family houses, the raw series is higher than the drop series in 2009–10 by up to 3.7 log points and lower in all eight quarters of 2011–12. The estimated path changes more than the cumulative decline: the peak-to-trough contrast is -0.64 log points (CI -3.03 to $+1.25$), which does not establish that the decline is unchanged. The path difference reaches 10.9 log points in the Bronx and is economically negligible for condominiums. Re-pairing A with B across the foreclosure (“bridge”) is a sensitivity to the pairing rule, not an estimate of recording error. Published New York indices fall far less than either ACRIS series and cannot discriminate between raw and drop (Section 4.4). The result survives narrower definitions of the dropped set: removing only referee’s deeds to lenders that carry cents and no securitization-type grantee, a third of the flagged deeds, still moves the 2010Q3 index by 2.6 log points (Section 4.5). Section 4.6 adds inference for the whole path, a placebo that removes ordinary pairs with the same timing, and explicit assignments of the unresolved review cases. Appendix C repeats these checks on pairs whose two endpoints are both 1–3 family deeds, with a band that holds jointly across six definitions. A hedonic index on the same deeds (Section 4.7) puts the raw index 2.07 log points in 2010Q3 above the index without them but shows no shortfall in 2011–12. The contrast is consistent with the influence of pairing a lender acquisition with its resale; it does not identify the market value of either transfer.

The closest literature studies distressed sales as *prices*: foreclosure discounts (Pennington-Cross 2006; Clauretie and Daneshvary 2009; Campbell et al. 2011; Conklin et al. 2023), externalities (Harding et al. 2009; Gerardi et al. 2015; Anenberg and Kung 2014; Fisher et al. 2015), and the sensitivity of house-price indices to including distressed transactions (Doerner and Leventis 2015). Index producers already screen for part of the problem: the S&P Cotality Case-Shiller methodology excludes, where the deed record identifies them, “reposessions of properties by mortgage lenders at the beginning of foreclosure proceedings”, while keeping subsequent sales by lenders as arm’s-length (S&P Dow Jones Indices 2025). The paper is closest to work showing that a recorded price can differ from the negotiated price for an institutional reason: Ben-David (2011) finds that financially constrained buyers inflated recorded prices to draw larger mortgages during the boom, and Nowak and Smith (2020) show that crisis-era repeat-sales indices are biased when the quality of the paired homes changes. Our point is prior to those: the legal basis for some registry amounts differs from a negotiated sale price, and a researcher building an index from public records inherits these amounts unless the deed is classified by its parties. A document review of a probability sample of 400 deeds, read by AI, measures how often the classifier agrees with those readings (Section 5).

2. Institutions: what the recorded amount means

2.1 Consideration under New York transfer-tax law

Every conveyance of real property in New York City is accompanied by state (TP-584 or TP-584-NYC) and city (NYC-RPT) transfer-tax filings. ACRIS, operated by the Office of the City Register within the New York City Department of Finance, publishes a document amount for each deed. ACRIS connects deed cover pages to transfer-tax filings; the public documentation reviewed here does not establish which form field supplies each published document amount. Section 3 tests whether those amounts behave as the tax definitions imply. State and municipal law both include discharged indebtedness in consideration, but their detailed foreclosure calculations differ. New York Tax Law §1401(d) and New York City Administrative Code §11-2101(9) include the cancellation or discharge of an indebtedness or obligation. New York Tax Law §1401(e) and the Rules of the City of New York (19 RCNY § 23-03) cover transfers in lieu of foreclosure and referee transfers. Schedule E of state form TP-584.1 (New York State Department of Taxation and Finance 2019) and 19 RCNY § 23-03(d) specify different calculations (City of New York 2024):

- **State foreclosure calculation (Schedule E, Part 1).** When the grantee is the mortgagee, its agent or nominee, or an entity it wholly owns, the starting amount is the foreclosure judgment — including interest accrued to the date of transfer — or the bid price, whichever is higher, plus continuing liens. For an unrelated buyer, it is the bid plus continuing liens.
- **City foreclosure calculation (19 RCNY § 23-03(d)(2)).** A referee’s deed is taxed on the bid, senior liens not canceled by the sale, and specified purchaser-paid costs, regardless of whether the buyer is the mortgagee.
- **State deed in lieu calculation (Schedule E, Part 2).** The starting amount is the unpaid secured debt, plus continuing liens and other amounts paid by the grantee.
- **City deed in lieu calculation (19 RCNY § 23-03(d)(1)).** Tax is computed on outstanding mortgage debt and unpaid accrued interest, whether or not the mortgagor is personally liable for that debt.
- **State recourse debt (Schedule E, Parts 1 and 2).** For a qualifying recourse debt, the state form uses the lower of the calculated amount and the property’s fair market value. This state-form cap is not stated in the cited city rule.

These legal calculations can produce an amount tied to the bid, judgment or debt rather than to a negotiated exchange price. The particular amount in an ACRIS deed record cannot be attributed to one of the two tax formulas from the registry field alone.

The state recourse cap limits any attempt to identify the exact basis of a recorded amount. New York permits deficiency judgments on residential mortgages, so some loans may qualify as recourse under the state form. Registry filings alone do not show whether a filer applied that cap, or how a subjective fair-market-value appraisal was made. The city rule cited above does not condition its debt-in-lieu calculation on personal liability. Digit patterns are therefore only an indirect empirical diagnostic: 74% of referee’s deeds to lenders record an amount with non-zero cents, characteristic of a debt or judgment balance rather than a rounded valuation, and their amounts sit well above subsequent resale prices. The state form cited is the July 2019 revision; we have not verified that earlier revisions computed consideration identically. We do not treat any recorded amount as proof of legal non-compliance or assume it equals debt. We test whether the class of amounts behaves like prices, and measure what including them does to an index.

2.2 Judicial foreclosure in New York

New York forecloses judicially. The court enters a judgment of foreclosure and sale and appoints a referee, who auctions the property and conveys it by referee’s deed to the successful bidder. When no third party outbids the plaintiff, the plaintiff — or its designee, servicer or securitization trustee — takes title. The lender then markets the property as real-estate-owned (REO) and sells it, usually within one to two years. In our data the median interval between a legal-rule deed and the next sale of the same parcel is 498 days.

Foreclosure timelines in judicial states such as New York are among the longest in the country (Cordell et al. 2015). New York added two procedural requirements during the crisis. Mandatory settlement conferences (CPLR 3408), introduced in 2008 for subprime and high-cost loans, were extended in December 2009 to all residential foreclosures on one-to-four family homes. Administrative Order AO/548/10 of 20 October 2010 required plaintiff’s counsel in residential foreclosures to affirm that the papers and notarisations had been checked for accuracy; AO/431/11 replaced it in March 2011, retroactive to November 2010. Both lengthened pending cases, and the timing appears directly in the registry (Section 3.3).

2.3 What ACRIS does and does not record

ACRIS publishes linked tables through NYC Open Data: a master record per document (type, dates, amount, percentage conveyed), one row per parcel conveyed, one row per party, cross-references to earlier documents, occasional free-text remarks entered at recording, and code lists. **No document type distinguishes a referee’s deed or a deed in lieu from an ordinary sale:** all are recorded as DEED. In the structured extract, the event must be inferred from its parties; the deed image itself can identify the instrument. A referee’s deed names the referee as grantor (“SMITH, ESQ., AS REFEREE”); a lender acquisition names a bank, servicer, trustee or federal agency as grantee. Staten Island deeds are recorded by the Richmond County Clerk and are absent from ACRIS; the analysis covers the other four boroughs. Cooperative apartments are conveyed by share transfer, not deed, and are outside the scope.

Table 1 sets out the taxonomy the rest of the paper uses. It separates the event (who conveys to whom) from the meaning of the amount. We treat three events as likely to carry legally defined consideration: a referee’s deed to a lender, any other conveyance to a lender (principally deeds in lieu), and transfers between lenders or to a federal agency. Third-party purchases at foreclosure auctions are competitive bids and are reported separately. Lender resales are market transactions and remain in the target population.

Table 1: Deed events and the meaning of the recorded amount (doc type DEED, four boroughs, 2003–2025)

Event	Parties	Amount	Deeds	Amount ≥ \$10k	Median (≥ \$10k)
Referee’s deed to lender	referee → lender or agency	bid, legal rule	6,742	77.3%	\$455,726
Other transfer to lender	owner → lender (deed in lieu etc.)	debt, legal rule	7,119	75.7%	\$477,067
Lender to lender	lender → lender or agency	legal rule	2,207	10.6%	\$283,436
Third-party auction	referee → non-lender	competitive bid	7,882	88.5%	\$400,530
REO resale	lender → non-lender	negotiated price	14,272	94.0%	\$321,000
Ordinary	neither party a lender or referee	negotiated price	1,354,845	62.5%	\$658,000

3. Data, classification and forensic evidence

3.1 Data

We retrieve the ACRIS master, parcel, party, cross-reference and remark tables and the DOF annualized sales and assessment rolls from the NYC Open Data API. The window is documents with identifiers from January 2003 to December 2025, which is when positive amounts begin to appear systematically. Socrata truncates responses silently at the requested limit, and the parcel table contains more than one row per document, so we fetch in calendar-month chunks. Each chunk is asserted equal to an exact count (1) taken immediately before and after the request, and the chunk totals are reconciled against a single live count for the whole window. All seven sources reconcile with zero drift: 3,579,646 deed and mortgage documents, 12,289,847 parcel rows, 24,885,207 party rows, 8,007,274 cross-reference rows, 217,958 remark rows, 845,607 DOF sales and 8,600,667 assessment records.

The deed file contains 1,470,324 deed-family documents in the four boroughs, of which 1,393,082 are of type DEED. The percentage-conveyed field is blank (recorded as 0) for 97.7% of 2003 deeds and 39.6% of 2004 deeds and almost never afterwards; those deeds look like ordinary sales (74% are multiples of \$1,000), so we treat 0 as unrecorded and drop only partial interests (between 0 and 100%). A buyer is coded as financed if a mortgage on the same parcel is dated between 7 days before and 30 days after the deed. Property class comes from the ACRIS property-type code: 1–3 family houses, 4–6 family buildings, residential condominium units and apartment buildings.

Generative AI tools (OpenAI’s ChatGPT and Codex, Anthropic’s Claude and Google’s Gemini) were used to draft and revise the text and to write and check the code, and the author reviewed their output. Section 5 describes the AI document review, and the Declarations give the full statement.

3.2 Classification

We classify each deed from its grantor and grantee names. A name is a lender if it contains an explicit institutional phrase from four families:

- federal agencies (Fannie Mae, Freddie Mac, the Secretary of Housing and Urban Development);
- banks and servicers (“BANK”, “SAVINGS”, “NATIONAL ASSOCIATION”, “MORTGAGE”, “SERVICING”, “HOME LOANS”);
- securitization vehicles (“CERTIFICATEHOLDERS”, “PASS-THROUGH”, “REO”);
- a short list of named servicers and bank REO subsidiaries.

A grantor is a referee if the name contains “REFEREE” or “SHERIFF”. A lien-holder that takes title on a referee’s deed (a New York City tax-lien trust or a condominium board) is treated as a lender. The rule is in Appendix A.

The design principle is that no ambiguous single word may act as a token. In the pilot, a bare “N.A.” token matched 19 Chinese and Korean surnames (“CHEN, NA”). At full scale the same risk appears in other forms:

- “LOAN” is a Vietnamese given name (“NGUYEN, LOAN”).
- “CHASE”, “SANTANDER”, “REO” and “BANK” are surnames.
- “GREENPOINT” is a Brooklyn neighbourhood that appears in developers’ LLC names.

The rule uses these words only inside institutional phrases. A name written “SURNAME, GIVEN” whose surname is itself a token is read as a person unless a corporate marker is present. We printed and read

every family of matches: the most frequent names in each category, random samples of hits, the names the person guard rejected, and the most frequent unmatched grantees. We then used the cents signature (Section 3.4) as a search tool for missed lenders. Among referee’s deeds to apparently non-lender buyers, amounts with cents pointed to lender REO subsidiaries the first rule missed: HomeSales Inc., Liquidation Properties, Property Asset Management, Retained Realty and Residential Funding. They were added. The audit and the list of changes between versions are part of the replication package.

Of 1,694,771 distinct grantee names on deed-family documents, 4,396 match the narrow rule. The classification labels 16,655 deed-family documents (16,068 of type DEED) as carrying a legal-rule amount.

3.3 Incidence

Figure 1 shows the count and share of legal-rule deeds with an amount of at least \$10,000. Their share of priced deeds was 0.5% in 2005–06, rose to 3.7% in 2008 and 4.2% in 2009, and fell back below 1% in 2011–12. The collapse coincides with the October 2010 attorney-affirmation requirement (AO/548/10). Referee take-backs fell from 959 in 2009 to 110 in 2011 and 81 in 2012. A second, smaller wave of deeds in lieu and post-foreclosure transfers followed in 2016–19, when the share reached 1.7%. Over the whole window there are 10,838 priced legal-rule deeds among 877,751 priced deeds (1.2%). By borough, 5,865 are in Queens, 2,706 in Brooklyn, 1,839 in the Bronx and 428 in Manhattan; by class, 9,466 are 1–3 family houses and 698 are condominium units.




Figure 1: Legal-rule deeds by year (doc type DEED, amount ≥ \$10,000). Top: counts by event. Bottom: share of all deeds with an amount of at least \$10,000.




Figure 2: Digit patterns of recorded amounts, screened sales in repeat pairs.

3.4 The amounts do not behave like prices

Negotiated prices are often rounded, while debt payoffs and judgment-based bids can accumulate interest, fees and costs to the cent. The implemented classifier does not use consideration, but cents patterns helped identify lender subsidiaries during rule development (Section 3.2). The digit distribution is therefore a diagnostic on the development data, not an independent validation, and cannot classify an individual deed (Figure 2 and Table 2). Among screened sales that enter a repeat pair:

- 69.9% of legal-rule amounts have non-zero cents, against 6.3% of ordinary prices.
- 17.5% are multiples of \$1,000, against 81.2%.

Referee's deeds to lenders are the least round of all: 74.1% carry cents and 14.1% are multiples of \$1,000. Third-party auction bids and REO resales sit between the two groups. The cents that do appear among ordinary sales are concentrated in condominium sponsor sales, so cents alone are a signal, not a label.

Subsequent resale paths provide a second diagnostic. Pairs whose first sale is a legal-rule deed fell in nominal terms 73.9% of the time (89.9% when the take-title date is in 2008–12), against 18.2% of pairs with no lender on either side. The median amount falls from \$449,490 at take-title to \$336,648 at resale. The lender acquires without a new mortgage in 6,845 of 6,883 such pairs and resells to a financed buyer in 3,592. This matches the pilot, in which cash-to-financed pairs outnumbered financed-to-cash pairs 143 to 1 (seed/).

Table 2: Forensic tests (screened DEED sales in repeat pairs; 2003–2025)

	Legal-rule	Referee's deed to lender	Other to lender	Third-party auction	REO resale	Ordinary
Sales in repeat pairs	9,091	4,624	4,312	4,096	10,054	387,872
Amount has cents	69.9%	74.1%	65.6%	14.3%	8.5%	6.3%
Multiple of \$1,000	17.5%	14.1%	21.0%	64.9%	49.5%	81.2%
Multiple of \$10,000	7.5%	5.3%	9.8%	32.5%	21.2%	44.3%
Pairs starting here: nominal fall	73.9%					18.2%

Notes: “Legal-rule” pools the three events of Table 1 whose amount is set by a legal rule. The last row compares pairs whose first sale is a legal-rule deed with pairs that have no lender or referee on either side.

3.5 External benchmarks

The class definitions and amount patterns motivate screening these transfers; they do not verify the pricing basis of every flagged deed. To assess where recorded amounts sit relative to other measures, we compare each screened 1–3 family sale with three benchmarks (Figure 3), each expressed as a median log gap $\times 100$.

DOF market value. For sales in 2010–18 and 2022–25 we match the final-roll market value for tax class 1 with a valuation date in January of the sale year; 222,849 of 226,687 eligible sales match. We subtract the median log ratio of ordinary sales in the same borough and year. Referee’s deeds to lenders sit 4.5 points above the ordinary benchmark: 13.1 points above in 2010–12, 1.1 above in 2013–18 and 13.3 below in 2022–25. Other lender acquisitions sit 1.5 below overall and 7.5 above in 2010–12. REO resales sit 33.7 points below, and third-party auction purchases 26.2 below. The rolls are matched on their January valuation date, which the extract dates of the 2023–2026 rolls confirm. Shifting the match a year either way leaves REO resales 30–36 points below and other lender acquisitions within 2 points of the benchmark, but moves referee’s deeds between -4 and $+8$, so only the small size of their gap, not its sign, is robust.

The parcel’s previous sale. We carry the previous ordinary sale of the same parcel forward with the clean index of Section 4. Legal-rule amounts sit at that value: the median gap is $+1.4$ for referee’s deeds and $+1.3$ for other acquisitions, compared with $+7.9$ for ordinary sales measured the same way.

The parcel’s next sale. We carry the next sale back to the take-title date, again with the clean index. The next sale is the lender’s resale 94% of the time. Here legal-rule amounts sit 34.2 (referee) and 23.9 (other) points *above*, while ordinary sales sit 11.0 points below their next sale.

The pattern is coherent. A legal-rule amount looks like the price a typical property of that kind would fetch, and like the parcel’s previous price carried forward — which is what a debt originated against that previous price should look like. It is well above what this particular property fetches when it next trades, by 35–45 log points relative to ordinary sales. Part of that gap is the forced-sale discount and the condition of a foreclosed house (Campbell et al. 2011), which the registry cannot separate from the recording rule. These comparisons cannot recover a take-title market value for a foreclosed house. They show that the recorded amount differs sharply from the next realised sale price while remaining too ordinary-looking to be caught by a plausibility screen.

3.6 Conventional screens keep them

Table 3 applies the screens common in repeat-sales work one at a time. A positive amount removes 13% of legal-rule deeds; a \$10,000 minimum removes a further 22%, mostly lender-to-lender transfers recorded at nominal amounts. Full-interest, single-parcel and residential-class filters remove almost none. At the end, 62.2% of legal-rule deeds survive, against 51.6% of ordinary deeds, and 54.6% enter a repeat pair, against 27.1%. Part of that gap in pairing comes from the screens themselves, which remove many zero-amount ordinary deeds. The rest comes from sales on either side of the deed: among deeds that pass every screen, 87.7% of legal-rule deeds and 52.5% of ordinary deeds enter a pair. The screens that most papers apply are designed for gifts and nominal transfers. A debt balance is neither.

The DOF annualized sales file, the standard public source for New York City transaction prices (e.g., Schwartz et al. 2003; Ellen et al. 2007), carries the same amounts. Of screened legal-rule deeds recorded in 2016–25, 93.8% of referee’s deeds to lenders and 95.3% of other lender acquisitions appear in the DOF file for the same parcel within seven days at the same amount (to the dollar). The rate for ordinary deeds




Figure 3: Recorded amount against three external benchmarks, 1–3 family houses. Median log gap $\times 100$. Left: relative to DOF market value, net of the median for ordinary sales in the same borough and year. Centre and right: relative to the value implied by the parcel's previous ordinary sale and its next sale, moved to the recording date with the clean (drop) index.

is 98.9%. The DOF file labels its parcel identifier with the condominium billing lot, so matching must use the borough, block and lot fields.

Table 3: Share of deeds surviving cumulative screens (deed-family documents, four boroughs, 2003–2025)

Screen (cumulative)	Legal-rule	Share	Ordinary	Share
All deed-family documents	16,655	100.0%	1,430,816	100.0%
Doc type DEED	16,068	96.5%	1,354,845	94.7%
Amount > 0	14,424	86.6%	872,576	61.0%
Amount $\geq$ \$10,000	10,838	65.1%	846,528	59.2%
Full interest conveyed	10,832	65.0%	840,379	58.7%
One parcel, no easement/air/partial lot	10,667	64.0%	795,142	55.6%
Residential class	10,361	62.2%	738,334	51.6%
In a screened repeat-sales pair	9,091	54.6%	387,872	27.1%

4. Consequences for a repeat-sales index

4.1 Estimator and corrections

We estimate the quarterly index of Bailey et al. (1963) by ordinary least squares,

$$\log R_{ib} - \log R_{ia} = \sum_{t=1}^T \delta_t (\mathbb{1}[t = b] - \mathbb{1}[t = a]) + u_q, \quad \delta_{2003Q1} = 0,$$

where R is the recorded amount and $q = (i, a, b)$ is a pair of consecutive screened sales of parcel i . Deeds pass the screens of Table 3. Pairs require a holding period of at least 180 days and an annualised log change no larger than one in absolute value. These screens were fixed before any full-population result

was computed. The screened sample has 762,062 sales and 238,290 pairs, of which 140,119 are 1–3 family houses and 86,943 are condominium units.

If a recorded amount departs from the corresponding price by $e_{ik} = \log R_{ik} - \log P_{ik}$, the dependent variable becomes $y_q^R = y_q^P + e_{ib} - e_{ia}$. A legal-rule deed L can enter up to two adjacent screened pairs with opposite signs. It enters the pair from the previous market sale, $A \rightarrow L$, as $+e_L$, and the pair to the lender’s resale, $L \rightarrow B$, as $-e_L$. If $e_L > 0$ – the amount exceeds what the property would fetch – the first pair overstates appreciation into the take-title quarter and the second overstates depreciation into the resale quarter. The net effect on δ_t depends on how these pairs are distributed over time and connected to the rest of the sample. No sign follows in general.

We report three corrections:

- **drop** removes every pair with a legal-rule endpoint and keeps the remaining pairs as they were.
- **bridge** removes legal-rule deeds from each parcel’s sale sequence *before* pairing, so that A and B form a new pair. This is how a researcher who filters deeds rather than pairs would proceed. The new pair spans the foreclosure: it carries the REO discount at B and any deterioration in condition, so bridge mixes the recording correction with a change in what is measured.
- **legs** removes only the $A \rightarrow L$ pairs, or only the $L \rightarrow B$ pairs, to show how each leg pushes the index.

We use drop to isolate the effect of excluding legal-rule endpoints under a fixed pairing rule. Dropping also changes which properties contribute to the index; it does not recover unobserved market prices for the lender take-title events. We report bridge as the sensitivity of the index to the pairing rule. Uncertainty comes from a property-clustered bootstrap: Poisson(1) weights per parcel, 300 draws, with the same draws for every variant, so that differences are paired. The largest of 92 quarterly differences is inflated by construction. We therefore also run a placebo (Section 4.6 reports a stricter, pair-level version): flag random *ordinary* deeds matched to the legal-rule deeds on borough $\times$ quarter counts, apply the drop correction, and record the largest quarterly difference; there are 200 draws.

4.2 Results

Figure 4 shows the raw and corrected indices for 1–3 family houses and their difference; Table 4 summarises all series. The largest quarterly difference is 3.69 log points in 2010Q3 (95% CI 2.48–4.92). This interval is pointwise at the selected quarter, not an interval for the maximum over all quarters. In the placebo, removing the same number of ordinary deeds at the same times produces a median maximum difference of 0.94 and a 95th percentile of 1.34. The largest placebo difference in 200 draws is 1.94.

The timing of the difference matters most. The raw index lies above the corrected one in 2008–10, with a pointwise band excluding zero in five quarters of 2009–10, by 1.1 points on average over 2008–10. It lies below in every quarter of 2011–12, by 2.3 points on average. The difference is often about one point negative, with its pointwise 95% band entirely below zero in 35 of the 52 quarters from 2013 to 2025, as later waves of lender acquisitions keep deflating resale quarters. Measured from annual averages, the raw index shows a decline of 30.6 points from 2006 to 2009 against 33.7 corrected, and a further decline of 12.6 points from 2009 to 2011 against 8.4 corrected. Over 2006–2011 the two agree within about a point (43.2 against 42.1). The peak-to-trough contrast, a decision criterion fixed before any full-population result was computed, agrees, though imprecisely (the written criterion did not fix the peak and trough windows used in Table 4): the raw decline is 49.4 log points, the corrected decline 50.1, and the difference is -0.6 (CI -3.0 to $+1.2$).




Figure 4: Top: quarterly repeat-sales index for 1–3 family houses in four boroughs, raw and with pairs that touch a legal-rule deed dropped. Bottom: raw minus clean, log points; solid with pointwise 95% bootstrap band: drop; dashed: bridge.

Figure 5 decomposes the difference into its two legs. Removing only the $A \rightarrow L$ pairs lowers the index at take-title dates, 2008–10, relative to purchase dates, 2005–07. Removing only the $L \rightarrow B$ pairs raises it at resale dates, 2011–12, relative to take-title dates. Each leg alone moves the index by 2–3 points at its peak; together they largely offset in the cumulative decline but not in its path. The sign of the $A \rightarrow L$ leg also matters for interpretation. Dropping pairs removes a distinctive group of properties as well as the legal-rule amounts, and if foreclosed houses lost market value faster than others, removing their pairs would *raise* the index at take-title dates. The $A \rightarrow L$ pairs instead push the index up, which fits an amount above market value more naturally than selection on distressed properties. This argument rests on the sign, not the magnitude, and does not separate the two channels.

The trough date is uncertain. The raw point estimate places the minimum in 2012Q1 and the corrected one in 2010Q3. But both series stay near their lows from late 2010 to early 2013, and across bootstrap draws the lag has a median of one quarter and a 90% range of –3 to +10 quarters. We make no claim about the trough date; the difference is in the path into and out of it.

The difference is concentrated where lender acquisitions are (Figure 6). For Bronx houses, where 11% of pairs touch a legal-rule deed, the raw index is 10.9 log points above the drop series in 2010Q3 (CI 6.6–14.7). Queens houses follow the citywide pattern, with a maximum of 4.0 (CI 2.2–5.6). In Brooklyn the dominant effect is a level shift after 2011 (maximum –3.1, CI –4.5 to –1.6). For condominium units, where legal-rule deeds touch 1% of pairs, the maximum is 0.8 log points. That exceeds every placebo draw (whose




Figure 5: Contribution of each leg to the raw 1-3 family index (raw minus index with that leg's pairs removed), log points. Dotted: both legs removed (drop).

maximum is 0.79), but its pointwise interval includes zero and it is economically negligible. Series for 4-6 family and apartment buildings have about 60 pairs per quarter and are too noisy to support conclusions either way.




Figure 6: Raw minus clean (drop), by borough and property type, with pointwise 95% bootstrap bands. Common vertical scale.

Table 4: Index error: raw minus corrected, log points (quarterly BMN index, 2003Q1–2025Q4)

Series	Pairs	Largest quarterly difference, drop [95% CI]	Quarter	Largest, bridge	Peak-to-trough diff., drop [95% CI]	Peak-to-trough diff., bridge [95% CI]
All residential	238,290	-2.37 [-2.92, -1.86]	2011Q1	+5.02	+0.56 [-0.58, +1.64]	-3.25 [-4.87, -2.12]
1–3 family	140,119	+3.69 [+2.48, +4.92]	2010Q3	+8.09	-0.64 [-3.03, +1.25]	-5.07 [-7.28, -2.11]
Condominium units	86,943	-0.82 [-2.29, +0.14]	2020Q1	-0.57	+0.20 [-0.20, +0.39]	+0.04 [-0.32, +0.19]
Bronx, 1–3 family	18,925	+10.92 [+6.64, +14.70]	2010Q3	+14.09	-2.43 [-11.48, +0.68]	-6.30 [-13.81, -2.73]
Queens, 1–3 family	73,010	+4.01 [+2.21, +5.63]	2010Q3	+9.43	-0.15 [-2.56, +2.58]	-5.83 [-7.97, -2.75]
Brooklyn, 1–3 family	46,348	-3.06 [-4.48, -1.57]	2011Q3	+5.23	+0.06 [-2.70, +1.56]	-1.06 [-5.66, +0.34]
Manhattan, condominium units	50,750	-0.93 [-3.16, +0.30]	2020Q1	-0.78	-0.33 [-0.74, +0.19]	-0.33 [-0.76, +0.09]

Notes: Quarterly confidence intervals are pointwise at the selected quarter; they are not simultaneous intervals for the maximum. Peak is the maximum over 2005–08 and trough the minimum over 2008–13. Placebo for 1–3 family (200 draws of ordinary deeds with matched borough $\times$ quarter counts): median of the largest difference 0.94, 95th percentile 1.34, maximum 1.94. For condominium units: 0.29, 0.51, 0.79.

4.3 The bridge correction

Bridging deserves separate comment because it is what a deed-level filter does by default. Replacing $A \rightarrow L$ and $L \rightarrow B$ with $A \rightarrow B$ lowers the corrected index further in 2009–10. The raw-minus-bridge difference reaches 8.1 log points for 1–3 family houses. The corrected peak-to-trough decline becomes 5.1 points larger than the raw one (CI 2.1–7.3). This is a pairing sensitivity, not an estimate of recording error: the added pairs span the foreclosure, almost all end in an REO resale, which sits 33.7 log points below the DOF market benchmark (Section 3.5) and, measured as in that section, 35.5 and 45.8 points below the values implied by the parcel’s previous and next sales, and they absorb any change in condition during default. The bridged index therefore includes the depreciation of foreclosed houses, which the drop index excludes. Either may be the object of interest, but the choice moves the crisis decline by five points and should be made explicitly.

4.4 An outside benchmark

Published indices give an outside check, although neither covers the same population (Figure 7). The S&P Cotality Case–Shiller New York index is a repeat-sales index of single-family homes across the metropolitan area. Its methodology already excludes lender repossessions that the deed record identifies (S&P Dow Jones Indices 2025). The FHFA all-transactions index for the New York–Jersey City–White Plains metropolitan division adds appraisals and covers only mortgages bought or guaranteed by Fannie Mae and Freddie Mac (U.S. Federal Housing Finance Agency, n.d.). We take both from FRED (Federal Reserve Bank of St. Louis 2026), convert them to quarterly log points and set them to zero in 2003Q1.

Two facts stand out. First, the ACRIS series rise and fall much further than either benchmark. Over the windows used in Table 4, Case–Shiller falls 30.4 log points from peak to trough and FHFA 20.1, against 49.4 for the raw series and 50.1 for drop. Annual FHFA indices for Bronx, Kings and Queens counties fall between 11.5 and 21.4 points, against between 40.4 and 48.9 for our borough series. The benchmarks cover wider areas and, for FHFA, a narrower set of loans, so the gap does not by itself show that any series is wrong. It is, however, far larger than the raw-minus-drop contrast, and none of these series should




Figure 7: ACRIS 1–3 family indices (raw and drop) against published New York indices, log points, 2003Q1 = 0. Case-Shiller New York (NYXRNSA) is monthly, averaged to quarters; FHFA all-transactions (ATNHPIUS35614Q) is quarterly. Both retrieved from FRED on 25 September 2026.

be read as the true price path of the city’s 1–3 family houses. Second, the benchmarks cannot choose between raw and drop. Over 2008–13, after removing the difference in level, drop lies slightly closer to both benchmarks: the root-mean-square gap is 5.6 against 5.9 for Case-Shiller and 6.6 against 6.9 for FHFA. Quarter-to-quarter changes in the raw series, however, track them more closely: the correlation is 0.71 against 0.57 for Case-Shiller and 0.22 against 0.14 for FHFA. Both benchmarks put the trough in 2012 (Case-Shiller 2012Q1, FHFA 2012Q2), as the raw series does. Given the flat bottom and the bootstrap spread in Section 4.2, we do not treat this as evidence about the trough date. The case for dropping legal-rule pairs therefore rests on the evidence in Section 3 that their amounts are not negotiated prices, not on agreement with a published index.

4.5 Narrower definitions of the non-market set

The document review (Section 5) leaves some flagged deeds unresolved, including referee’s deeds to grantees named like REO investors. To check that the index result does not rest on such cases, we redo the drop correction for 1–3 family houses with narrower sets of dropped deeds, and with the classifier’s broad option as an upper side (Table 5). Every set comes from the same classifier:

- **D1** keeps only referee’s deeds to lenders.
- **D2** removes deeds whose grantee falls in the securitization category of Appendix A, which holds the “REO” token. It removes five of the seven ambiguous REO-named cases in the review sample.
- **D3** keeps only amounts with non-zero cents.
- **D4** requires all three.
- **D5** uses the broad option of Section 5.

Each definition gets its own placebo, as in Table 4: 200 draws of random ordinary deeds that the definition does not flag, matched to its flagged deeds on borough $\times$ quarter counts.

Table 5: Raw minus drop for 1–3 family houses under alternative definitions of the dropped set, log points

Dropped set	Flagged 1–3 family deeds	Pairs dropped	Largest quarterly difference [95% CI]	Quarter	Placebo maximum	Peak-to-trough diff. [95% CI]
D0 published	9,330	12,483	+3.69 [+2.53, +4.93]	2010Q3	1.94	−0.64 [−3.06, +1.19]
D1 referee’s deeds only	4,719	6,824	+2.89 [+1.96, +4.01]	2010Q3	1.34	+0.81 [−1.16, +1.32]
D2 no securitization category	8,743	11,683	+3.61 [+2.58, +4.92]	2010Q3	1.78	−0.40 [−2.88, +0.93]
D3 non-zero cents	6,515	8,770	+3.64 [+2.66, +4.78]	2010Q3	1.47	−0.24 [−2.77, +0.72]
D4 narrowest (D1, D2 and D3)	3,309	4,754	+2.63 [+1.83, +3.67]	2010Q3	1.19	+0.63 [−0.77, +0.98]
D5 broad option	11,368	14,010	−3.63 [−4.60, −2.56]	2011Q3	1.66	−0.39 [−2.96, +1.41]

Notes: Intervals are pointwise, from a paired Poisson(1) parcel bootstrap with 300 draws. These draws differ from Table 4’s, so D0’s interval differs in the second decimal. Placebo maximum: the largest of 200 placebo draws of the largest quarterly difference. No placebo draw reaches the observed value under any definition.

Under every definition the largest difference exceeds all 200 of its placebo draws. Dropping only the narrowest set, 3,309 deeds or about a third of the published set, still moves the 2010Q3 index by 2.63 log points (pointwise 95% CI 1.83–3.67), against a placebo maximum of 1.19. Removing the securitization category changes the maximum from 3.69 to 3.61. Five definitions peak in 2010Q3. The broad option peaks in 2011Q3, with the raw index below; its 2010Q3 difference is +3.55. As in Table 4, the peak-to-trough interval includes zero under every definition: the quarterly path changes, and none of these checks establishes a change in the cumulative decline.

The narrower sets are not an independent validation. Non-zero cents do not prove a computed bid or debt, and an institutional name on a referee’s deed does not prove that the buyer is the creditor. D2 also keeps two ambiguous grantees that the rule classes as banks. A pair’s property class follows its first sale, so some dropped pairs touch a flagged deed of another class: 44 under D0 and 10 under D4. The placebo does not match those endpoints. Appendix C repeats the analysis on pairs whose two endpoints are both 1–3 family deeds.

4.6 Inference for the whole path, an aligned placebo and unresolved cases

Three further checks address how Tables 4 and 5 were built. Table 6 reports the first two and Table 7 the third.

An aligned placebo. The placebo in Tables 4 and 5 removes ordinary deeds. That need not remove as many pairs as the drop correction, or pairs with the same timing, and it leaves out flagged endpoints of other property classes. Here the placebo removes ordinary pairs instead. Each of 200 draws deletes exactly as many 1–3 family pairs with two ordinary endpoints as the correction drops, matched on borough, first-sale quarter and second-sale quarter. For D0, 91% of pairs match in their exact cell; the rest match on borough and sale years, then on sale years alone, then on the year of the second sale. It is a much harder test. Removing 12,483 random ordinary pairs with the same timing moves the index by up to 3.41 log points, against 1.94 in the deed placebo. The observed 3.69 still exceeds all 200 draws, but only just. The narrowest set clears it by more: 2.63 against a placebo maximum of 1.89. A whole-path statistic, the root-

mean-square difference over all quarters, exceeds every draw under both definitions (1.33 against at most 1.27 for D0, and 0.67 against 0.60 for D4).

Inference for the whole path. The intervals in Tables 4 and 5 are pointwise at the quarter where the difference is largest, which was chosen after seeing the series. A simultaneous 95% band covers all quarterly differences at once (sup-t, from the same paired parcel bootstrap). It is wider: at 2010Q3 it runs from 1.59 to 5.78 for D0 and from 0.88 to 4.37 for D4. It excludes zero in 28 quarters for D0 and 23 for D4, including 2010Q3 and every quarter from 2011Q3 to 2012Q3. Section 4.2 describes the path by window averages, which now have intervals: the raw index lies above the corrected one over 2008–10 by 1.07 log points (95% CI 0.55–1.59) and below it over 2011–12 by 2.28 (1.72–2.76).

Table 6: Whole-path inference and an aligned placebo, raw minus drop for 1–3 family houses, log points

	D0 published	D4 narrowest
Pairs dropped	12,483	4,754
Difference in 2010Q3	+3.69	+2.63
Simultaneous 95% band at 2010Q3	[+1.59, +5.78]	[+0.88, +4.37]
Quarters whose simultaneous band excludes zero	28	23
Mean difference 2008–10 [95% CI]	+1.07 [+0.55, +1.59]	+1.05 [+0.83, +1.24]
Mean difference 2011–12 [95% CI]	–2.28 [–2.76, –1.72]	–1.09 [–1.27, –0.91]
Mean difference 2013–25 [95% CI]	–1.00 [–1.43, –0.57]	–0.25 [–0.38, –0.14]
Aligned placebo, largest difference: median, 95th percentile, maximum	2.68, 3.11, 3.41	1.50, 1.77, 1.89
Root-mean-square difference, observed	1.33	0.67
Aligned placebo, root-mean-square: median, 95th percentile, maximum	0.87, 1.12, 1.27	0.45, 0.53, 0.60

Notes: Paired Poisson(1) parcel bootstrap with 300 draws, as in Table 5. The simultaneous band uses the 95th percentile of the largest standardised deviation across quarters (critical value 3.27 for D0, 3.66 for D4). The aligned placebo deletes random pairs with two ordinary endpoints, matched to the dropped pairs on borough $\times$ first-sale quarter $\times$ second-sale quarter, with coarser cells where a cell runs short; 200 draws. No placebo draw reaches the observed value of either statistic.

Unresolved classifications. The document review leaves deeds unresolved in every stratum (Section 5). Table 7 assigns each group against the classifier and re-estimates the index. R1 returns 10.7% of flagged deeds to the sample, chosen at random: the share of the flagged sample read as a market sale or left unresolved, which gives the lower precision bound. A treats every third-party auction purchase as non-market, since 48 of the 50 sampled are unresolved. S3 and S4 drop, at random, the share of the suspicious and random unflagged strata read as non-market or left unresolved (15% and 12%); most unresolved random deeds are possible family transfers. ALL combines the four. The strata are rebuilt from the data and match the frozen stratum sizes exactly.

Table 7: Raw minus drop for 1–3 family houses when unresolved review cases are assigned against the classifier, log points

Scenario	Pairs dropped	2010Q3	Mean 2008–10	Mean 2011–12	Peak-to-trough diff.
D0 as published	12,483	+3.69	+1.07	–2.28	–0.64
R1 flagged, 10.7% returned	11,165	+3.24 [+2.76, +3.62]	+0.91 [+0.78, +1.08]	–2.02 [–2.15, –1.86]	+0.03 [–0.38, +0.41]
A auction purchases dropped	16,701	+2.96	+0.82	–2.36	+0.41
S3 suspicious, 15% dropped	12,776	+3.68 [+3.50, +3.81]	+1.09 [+0.98, +1.15]	–2.29 [–2.40, –2.22]	–0.60 [–0.75, –0.45]
S4 random, 12% dropped	40,617	+3.72 [+2.14, +5.30]	+1.08 [+0.02, +2.28]	–2.33 [–3.25, –1.08]	–0.93 [–2.04, +0.61]
ALL four together	43,452	+2.63 [+0.47, +4.40]	+0.58 [–0.44, +1.76]	–2.02 [–3.17, –0.89]	+0.47 [–1.02, +1.65]

Notes: Scenarios with a random share report the median and, in brackets, the 2.5th and 97.5th percentiles over 200 draws; the brackets show which deeds are assigned, not sampling error. Pairs dropped: median over draws.

In every scenario the raw index lies above the corrected one in 2010Q3 and about two points below it over 2011–12. Returning flagged deeds shrinks the 2010Q3 difference from 3.69 to 3.24. Dropping auction purchases moves the largest difference to 2011Q3, at –3.37. S4 and ALL drop a random eighth of ordinary deeds, about 28,000 further pairs, which adds noise: under ALL the 2008–10 average ranges from –0.44 to 1.76 across draws, while the 2011–12 average stays between –3.17 and –0.89. The peak-to-trough difference stays between –2.0 and +1.7.

The path result survives these checks in a narrower form. Within the repeat-sales estimator, the raw index’s shortfall in 2011–12 holds up best: it clears the simultaneous band and appears in every classification scenario. The excess in 2009–10 clears the aligned placebo and the simultaneous band under the published and narrowest definitions, but it depends more on how unresolved deeds are assigned. On the stricter pair universe of Appendix C, a band that holds jointly across six definitions and all quarters excludes zero in five of the eight 2011–12 quarters for the published definition and in none of them for the narrowest. The narrowest definition’s 2011–12 average is negative, with a separate bootstrap interval that excludes zero; that interval is not joint across definitions or quarters.

4.7 A hedonic index

Sections 4.2–4.6 use one estimator. A hedonic index tests whether the path difference depends on it. It uses screened sales without pairing successive transfers. An observation can also affect fitted covariate coefficients and normalization, so its influence need not remain confined to its recording quarter.

We regress the log recorded amount of each screened 1–3 family sale on quarter, ZIP-code and building-class indicators, log gross and land area and their squares, building age and its square, indicators for two and three units, and the number of commercial units. Attribute prices are fixed over 2003–2025, and the index is the quarter coefficients, with 2003Q1 as the base. The attributes come from the Department of Finance’s annual sales files. Each deed is joined to the record of the same parcel with the nearest sale date within 45 days. Deeds with a tied match or none are left out, and the join does not require the two prices to agree, so the outcome cannot choose its own controls. Of 463,245 screened sales, 445,749 have a unique match, and 435,486 sales on 284,880 parcels have usable attributes and fall in ZIP and class cells with at

least 20 sales. The sample keeps 8,341 of the 9,330 flagged deeds in the published set (D0) and 2,983 of the 3,309 in the narrowest (D4, Section 4.5). In 433,219 of the 435,486 sales the two recorded prices agree within 0.5%.

The annual files for 2003–2010 state that their descriptive data are as of 6 May 2012, so size and age can describe the house years after the sale; building class is recorded at the time of sale. We therefore also fit a model with only the quarter, ZIP and class indicators, on the same rows. Uncertainty comes from 300 paired Poisson(1) parcel-bootstrap draws with full refits. The simultaneous band in Table 8 covers both models, both definitions and all 91 non-base quarters at once (critical value 3.96).

Table 8: Raw minus screened for 1–3 family houses: hedonic indices, and repeat sales on the hedonic sample, log points

	D0 published	D4 narrowest
Flagged deeds kept in the hedonic sample	8,341 of 9,330	2,983 of 3,309
Hedonic, all controls: 2010Q3 [simultaneous 95% band]	+2.07 [+1.27, +2.87]	+1.55 [+0.94, +2.16]
Hedonic, class and location only: 2010Q3 [simultaneous 95% band]	+1.89 [+1.12, +2.66]	+1.36 [+0.80, +1.92]
Hedonic, all controls: largest difference (quarter)	+2.21 (2009Q1)	+1.59 (2010Q1)
Hedonic, all controls: mean 2008–10 [95% CI]	+1.39 [+1.25, +1.53]	+1.05 [+0.98, +1.12]
Hedonic, all controls: mean 2011–12 [95% CI]	+0.14 [+0.02, +0.27]	+0.10 [+0.06, +0.13]
Hedonic, all controls: mean 2013–25 [95% CI]	+0.02 [−0.08, +0.14]	+0.05 [+0.02, +0.08]
Repeat sales, both endpoints in the hedonic sample: pairs dropped	10,480 of 124,768	3,980 of 124,768
Repeat sales, same sample: 2010Q3	+2.99	+2.34
Repeat sales, same sample: mean 2011–12	−2.31	−1.03

Notes: Hedonic indices: log recorded amount on quarter, ZIP-code and building-class indicators, with the structural attributes listed in the text (all controls) or without them (class and location only); 435,486 sales, 2003Q1 = 0. Simultaneous bands: one critical value covers both models, both definitions and all 91 non-base quarters. Window means: percentile intervals from the same 300 draws. Repeat sales: the drop correction of Section 4.1 on pairs whose two endpoints are both in the hedonic sample; point estimates only. On the full sample (Table 6) the 2011–12 means are −2.28 and −1.09.

Table 8 and Figure 8 give the results. In 2010Q3 the raw hedonic index lies 2.07 log points above the index without D0 deeds (simultaneous band 1.27–2.87) and 1.55 above the index without D4 deeds (0.94–2.16). Without the structural controls the differences are 1.89 and 1.36, and every band at 2010Q3 excludes zero. The excess covers 2008–10, at 1.39 log points on average for D0, and peaks at 2.21 in 2009Q1.

After 2010 the two estimators part. The hedonic difference averages 0.14 log points over 2011–12 for D0 and 0.10 for D4. The repeat-sales index on pairs whose two endpoints are both in the hedonic sample keeps its shortfall, −2.31 for D0 and −1.03 for D4, against −2.28 and −1.09 on the full sample (Table 6). The common-sample comparison rules out attribute coverage as the sole explanation. It is consistent with the pairing mechanism described in Section 4.1: the amount enters consecutive repeat-sale pairs with opposite signs (Figure 5). The estimators still impose different specifications and weights, and neither identifies the unobserved market price. After 2013 the repeat-sales difference remains near −1 while the hedonic difference averages 0.02.

The excess in 2009–10 therefore does not depend on the estimator. The shortfall in 2011–12 is specific to repeat-sales indices, which are the subject of this paper: a researcher who builds a hedonic index from




Figure 8: Raw minus screened hedonic index for 1–3 family houses, log points. Shaded: simultaneous 95% bands for the model with all controls, across both definitions and all quarters (critical value 3.68), narrower than the bands in Table 8, which also cover the class-and-location model. Dotted: the repeat-sales difference of Figure 4 (drop, published set).

these records would see the first and not the second. Neither model verifies housing quality, and the join is treated as fixed, so its errors are outside the bootstrap.

5. A filtering method for registry-based price research

The practical recommendation is short.

1. **Classify deeds by their parties, not their document type.** In ACRIS, referee’s deeds and deeds in lieu are recorded as ordinary DEEDs. The parties are in a separate table and must be joined.
2. **Remove deeds whose amount is set by a legal rule:** referee’s deeds to lenders, other conveyances to lenders, and lender-to-lender transfers. Keep lender resales, which are market sales. Report third-party auction purchases separately.
3. **Drop the pairs, and decide separately whether to bridge.** Dropping excludes legal-rule endpoints from a fixed set of pairs. Bridging adds new pairs across foreclosure and changes the population being measured.
4. **Audit the name rule on the data it will be applied to.** Print the most frequent matches in each category, a random sample of matches, the names a person-name guard rejects, and the most frequent unmatched grantees. Use the cents signature to search for missed lenders among referee’s-deed buyers.
5. **Do not rely on price-plausibility screens.** Legal-rule amounts sit near assessed values and near the prior sale carried forward. A ratio screen will not find them.

The classifier implementing steps 1, 2 and 4 is released as a single Python module, `acris_consideration_filter.py`. It has no dependencies beyond the standard library, plus pandas for the vectorised interface. It classifies each distinct name once and returns, for each document, the event label, whether the amount is market, and the name flags. A broad option adds generic finance words (FUNDING, FINANCE, CAPITAL) and series-code patterns. They add matches, but the broad-only matches include relocation companies and property investors as well as lenders, and their effect on precision and recall is unmeasured.

A documentary check. The digit signature (Table 2) and the external benchmarks (Section 3.5) show that the labelled class behaves as the mechanism predicts. They do not measure the classifier’s error rates. Two ACRIS tables give evidence that does not depend on party names. Filers sometimes explain a deed in a free-text remark (“THIS IS A REFEREE’S DEED. SEE TERMS OF FORECLOSURE SALE.”), and deeds sometimes cross-reference earlier documents.

- **Remarks.** 542 deeds carry a remark describing a foreclosure sale. The classifier labels 207 of them lender take-backs; an AI reading of the buyer names classes every flagged buyer as a bank, servicer, trustee, agency or lender REO vehicle. Among the 335 it does not flag, AI readings of the buyer names identify 13 more lenders or lien-holders (for example “TCIF REO1, LLC”, “BANKUNITED”, misspelt HUD entries and two condominium boards) and 13 possible ones. Within this convenience sample, the flagged share of name-identified lender buyers is therefore between 88.8% and 94.1%; this is not a recall estimate.
- **Deeds in lieu.** Of 21 deeds whose remarks say “deed in lieu”, 12 are flagged. The other nine are mostly commercial buildings conveyed to private lenders organised as LLCs, which no name-based rule can identify.
- **Cross-references.** Flagged deeds cite a mortgage far more often than other deeds: 7.3% of other transfers to lenders and 4.2% of referee’s deeds to lenders, against 0.03% of ordinary deeds.

Filers write remarks for their own administrative or legal reasons, so the sample is not random and cannot sign any population-level correction. It was selected independently of the names the rule reads and was not used to tune the rule. The name readings are AI assessments of names, not readings of the documents. The misses are listed in `audit/remark_check_reading.csv` and logged as candidates for a later version.

The document review. To measure error rates on a probability sample, we drew a stratified sample of 400 screened sales from the 411,113 that enter repeat pairs:

- 150 labelled legal-rule, to estimate precision;
- 50 third-party auction purchases;
- 100 unlabelled deeds with a suspicious signature (cents and a subsequent nominal fall, or a broad-only match), to find misses where they are likely;
- 100 other unlabelled deeds, drawn at random.

Each deed is linked to its ACRIS document image. For every deed the reading records the instrument, whether the grantee is a lender or an affiliate, the basis of the consideration and whether the transfer is a market sale, without the label or the stratum. The original protocol called for two human readers; this version uses AI instead. OpenAI’s ChatGPT read each deed image and cited the pages it relied on. Anthropic’s Claude then checked every reading against the city’s ACRIS open data, covering document type, amount and every party’s name and role, and against the protocol’s label rules, and sent back the readings it could not reconcile for a second look. Claude did not read the images, so the two steps are not independent readings. The error rates below measure agreement with these readings, not accuracy against a verified record. ACRIS does not allow automated capture from its website (New York City Department of Finance, Office of the City Register, n.d.), so the images were opened one at a time in an ordinary browser session and none were downloaded in bulk. Precision and recall are ratios of Horvitz–Thompson totals using the recorded inclusion probabilities, with approximate stratified-design intervals. A deed whose market-sale reading remains uncertain is reported as unresolved, never recoded.

What the review finds. Of the 400 deeds, 316 have a definite market-sale reading and 84 remain unresolved: 15 of the 150 flagged deeds, 48 of the 50 third-party auction purchases, 11 of the 100 suspicious

unflagged deeds and 10 of the 100 random unflagged deeds. After the second model's check, 7 readings were revised, all on the market-sale label; `audit/cross_ai_review/review_log.csv` lists each revision and its reason. Four referee's deeds to a documented lender, agency or trustee moved from uncertain to non-market, the rule the first model had applied to similar deeds elsewhere in the sample. Three non-market calls that the protocol does not support became uncertain.

Among resolved deeds, the classifier's precision is 0.993 (95% interval 0.978–1.000). One flagged deed in the sample is a market sale: a purchase by a family trust whose administrative trustee is a bank. Its recall is 0.496 (0.168–0.824). The protocol counts every transfer whose amount is not a negotiated price as non-market, including sales between relatives, and the classifier does not look for those. Two family transfers in the random stratum, each standing for about 3,900 deeds, carry most of the missed weight. Counting only foreclosure-type transfers (referee's deeds, deeds in lieu and other transfers to lenders) as the target, recall is 0.966 (0.934–0.998). The four misses in that class are three lender-side buyers without an institutional word in their names (a Novastar real-estate holding company, LNV Corporation and Radian Services) and an individual plaintiff who bought the property at the foreclosure sale.

If every unresolved deed went against the classifier, precision would be between 0.893 and 0.993 and recall between 0.135 and 0.523. The lower recall bound is low because each of the ten unresolved random deeds carries a large weight; most are sales with a possible family link or an unusual entry on the transfer form. Scoring the first model's readings without the revisions gives a precision and recall of 0.992 and 0.490.

The auction stratum is almost entirely unresolved for a structural reason. A referee's deed records the highest bid but never shows whether anyone bid against it, and the protocol accepts an auction purchase as a market sale only when the image shows a competitive sale. The deed image alone cannot settle whether a third-party auction price is a market price.

The analyst key that links each audit identifier to its stratum, label and inclusion probability is deposited with the replication package as `audit/audit_key.csv`. It was opened only after the readings and revisions were locked, and its SHA-256, `ed750a76dc10fb418d1dee4902bd15efc82816c7045f447e02857631effcaf0e`, matches the value published with an earlier public version of the package, before any reading began.

The review bears on three known limits of the classifier.

- **False negatives among auction buyers.** 14% of third-party auction amounts carry cents, well above the 6% for ordinary sales. The remark check suggests that some of those buyers are lender vehicles without an institutional word in their name. The review finds one such buyer among the 50 sampled auction purchases and two among the suspicious unflagged deeds.
- **Referees without the word “referee”.** Only 61% of remark-confirmed foreclosure sales name the grantor as a referee. The others give the attorney's name alone (“MELTZER, CRAIG ESQ.”). Lender take-backs among them are still flagged through the buyer, but the split between referee's deeds and other lender acquisitions in Table 1 understates referee's deeds. In the review sample, 25 deeds read as referee's deeds name the grantor in the index only as an attorney.
- **Ambiguous “REO” grantees.** The rule treats grantees such as “NYC REO LLC” as lenders; some may be investors buying at market. Grantees with “REO” in the name account for fewer than 600 priced deeds. The review leaves this open: seven referee's deeds to grantees such as “REO Management Corp.”, “RR REO II LLC” or a mortgage trust stay unresolved because the deed does not say whom the grantee acts for. Counting them as lender take-overs leaves precision at 0.993 and raises recall to 0.509. Section 4.5 shows that the index result does not depend on them.

6. Conclusion

Registry amounts arise from a recording and transfer-tax process. In New York, foreclosure and deed-in-lieu rules can base consideration on a bid, judgment or debt; the calculation behind each public ACRIS amount is not observed. The amounts on lender take-title deeds look like prices: they are near assessed values and the parcel's past price, survive the screens researchers use and appear in the city's own sales file. Their digit patterns and subsequent resales are consistent with their being unsuitable as market prices for these events.

The effect on a standard repeat-sales index is specific. It is concentrated in 1–3 family houses in the Bronx, Queens and Brooklyn, in and after the crisis. The citywide path changes: the raw index falls less into 2009–10 and more into 2011–12 than the drop series. A hedonic index shows the first movement without the later shortfall, consistent with the role of repeat-sale pairing. The point estimates of the cumulative peak-to-trough decline differ by only 0.6 log points, and the interval for that contrast includes zero. For Bronx houses the quarterly difference reaches eleven log points; for condominiums it is negligible. The trough date is uncertain. Bridging across foreclosure moves the decline more, by changing the pairs rather than correcting recorded amounts.

Once the problem is known, the correction is mechanical: join the party table, classify the deed and drop the pairs. Among the 316 resolved AI readings, estimated precision is 0.993 and foreclosure-type recall is 0.966; 84 sampled deeds remain unresolved. These rates condition on the resolved readings and do not establish independent accuracy. The index-path difference persists when only the narrowest flagged set is dropped. For the published definition, the largest difference exceeds every draw of a pair-level placebo, narrowly, and the raw index's shortfall in 2011–12 clears a simultaneous band and appears in each assignment scenario for the deeds the review leaves unresolved (Tables 6 and 7). Before registry amounts are used as prices, researchers should check which events record amounts under a legal rule and whether conventional screens remove them.

Appendix A. Classifier rule (module version 1.1.0)

Names are upper-cased, “&” is replaced by “AND”, periods are removed and other punctuation is replaced by spaces. A name is a *lender* if, after normalisation, it matches one of:

- **Agency:** FANNIE MAE; FEDERAL NATIONAL MORTGAGE; FREDDIE MAC; FEDERAL HOME LOAN MORTGAGE; GOVERNMENT NATIONAL MORTGAGE; SECRETARY OF (THE DEPARTMENT OF) HOUSING/HUD; HOUSING AND URBAN DEVELOPMENT; SECRETARY OF VETERANS; VETERANS AFFAIRS; FEDERAL DEPOSIT INSURANCE.
- **Bank or servicer:** BANK; BANCORP; BANKING; SAVINGS; NATIONAL ASSOCIATION; FSB; CREDIT UNION; MORTGAGE (including compounds such as CITIMORTGAGE); MTG; SERVICING; SERVICER; HOME LOAN(S); LOAN SERVIC...; LOAN TRUST; SAVINGS/INVESTMENT/BUILDING AND LOAN; LENDING; LOAN(S) followed by a corporate suffix.
- **Securitization:** CERTIFICATEHOLDERS; PASS-THROUGH; ASSET-BACKED; LEGAL TITLE TRUST; PARTICIPATION TRUST; REMIC; SECURITIZATION; REO.
- **Named servicers and REO subsidiaries:** a fixed list in the released module, including Nationstar, Ocwen, Selene Finance, HomeSales, Liquidation Properties, Property Asset Management, Retained Realty and Residential Funding. The code is the exact pattern specification.

Person guard. A match is rejected if the name has the form “SURNAME, GIVEN” and all of the following hold:

- the one-word surname equals the matched token;
- the given-name part has at most four words and no digits;
- the given-name part has no corporate or institutional word (INC, LLC, CORP, TRUST, NA, SUCCESSOR, MERGER, ...).

Referee. REFEREE, REFERE, REFREE, AS REF, SHERIFF.

Lien-holder. NYCTL, TAX LIEN, BOARD OF MANAGERS; counted as a lender only on a referee’s deed.

Event. Let g be “some grantee is a lender (or a lien-holder on a referee’s deed)”, r “some grantor is a lender” and f “some grantor is a referee”. The events are:

- $f \wedge g$: referee’s deed to lender;
- $f \wedge \neg g$: third-party auction;
- $\neg f \wedge g \wedge \neg r$: other transfer to lender;
- $\neg f \wedge g \wedge r$: lender to lender;
- $\neg f \wedge \neg g \wedge r$: lender resale;
- otherwise: ordinary.

Appendix B. Replication

`run_all.sh` fetches the public NYC Open Data files, builds the deed classifier, estimates the index variants, and reproduces the tables and figures. Stage 06 compares the indices with the published benchmarks stored in `data/external/`. Stage 05 checks the paper’s numbers against the derived data. Stage 07 (`07_narrower_definitions.py`) re-estimates the drop correction under the alternative definitions of Section 4.5; with `--check` it confirms Table 5 against its output. `score_cross_ai_review.py` scores the frozen audit sample from the AI readings and the logged revisions with the estimator of `deed_review.py`, after checking the analyst key’s hash; with `--check` it confirms the review numbers in Section 5. The README gives the commands for using the archived data snapshot.

The go/no-go criteria in `notes/G0_NO_GO.md` were written before any full-population result was computed. They are reported with the outcome, including the finding that the fixed two-point threshold was less conservative than the placebo.

The replication package, including the classifier module, the frozen audit sample and its key, is archived on Zenodo (doi:10.5281/zenodo.22925383). Stage 08 (`08_path_checks.py`) produces Tables 6 and 7, and stage 09 (`09_both_endpoint_checks.py`) produces Table 9 from the frozen deed file; with `--check`, each confirms the corresponding text against its output. Stage 10 (`10_hedonic_checks.py`) produces Table 8 and Figure 8 from hedonic estimates that join each deed to the Department of Finance’s annual sales files; with `--check` it confirms Section 4.7.

Appendix C. Pairs with two 1–3 family endpoints

A pair’s property class in Section 4 follows its first sale. Some 1–3 family pairs therefore end at a deed classed otherwise, and some dropped pairs touch a flagged deed of another class (Section 4.5). This appendix repeats the drop correction on a stricter universe: pairs whose two endpoint deeds are both classed as 1–3 family houses. That removes 1,739 of the 140,119 pairs in Tables 4–7 and leaves 138,380, with 237,570 distinct endpoint deeds. The screens, estimator and definitions D0–D5 are those of Section 4.5, and no new pairs are formed. The appendix uses its own 999 parcel-bootstrap draws and 999 placebo

draws per definition. Its numbers answer the same question on a different denominator, so they should not be substituted into Tables 4–7.

The path. The largest difference is again in 2010Q3: +3.61 log points for D0 (+3.69 in Table 5) and +2.62 for D4 (+2.63). D1, D2 and D3 also peak in 2010Q3, at +2.88, +3.53 and +3.59; the broad option D5 peaks in 2011Q3 at –3.70. Under both D0 and D4 the raw index lies above the drop series over 2008–10 by about one point on average. It lies below it in every quarter of 2011–12, by 2.28 points on average for D0 and 1.08 for D4 (Table 9). These averages are close to those of Table 6. The peak-to-trough interval includes zero under both definitions.

A band across definitions. Table 6’s simultaneous band covers all quarters for one definition at a time. Here we form a single 95% band that holds jointly over all six definitions and all 91 quarters after the 2003Q1 normalization. It is a centered sup-norm band from 999 common Poisson(1) parcel-weight draws. Unlike the studentized band of Table 6, it has the same half-width, 1.98 log points, at every quarter and definition, so the two bands are not comparable quarter by quarter. At 2010Q3 the joint band runs from +1.63 to +5.59 for D0 and from +0.64 to +4.60 for D4. For D0 it excludes zero in ten quarters, five of them in 2011–12 (2011Q1, 2011Q3, 2011Q4, 2012Q1 and 2012Q2). For D4 it excludes zero only in 2010Q1 and 2010Q3, and in none of the eight quarters of 2011–12: D4’s quarterly differences there are all negative, but none is larger in absolute value than the band’s half-width. A separate question is whether D4’s 2011–12 average differs from zero. That average, –1.08, has a 95% bootstrap interval of –1.30 to –0.88. The interval is for this one window and this one definition. It is not joint across definitions or quarters, and it does not change what the joint band shows. The joint band treats the classifier’s labels as given. It reflects resampling of parcels, not misclassification, and it does not indicate which definition is correct.

Two placebos for this universe. The first draws ordinary deeds from the universe’s endpoints, matched to each definition’s flagged endpoints on borough $\times$ quarter counts, and applies the drop correction. No cell runs short. In 999 draws the largest placebo difference is 2.59 for D0 and 2.06 for D4, and no draw reaches the observed value under any of the six definitions. This placebo deletes fewer pairs than the correction: for D4, 3,555 on average against 4,700. A random ordinary deed less often sits between two eligible pairs than a flagged deed does. The second placebo, run for D4, therefore also matches each deed on the number of eligible pairs it enters, one or two. It deletes 4,664 pairs on average (range 4,645–4,680). Its 95th percentile is 1.74, and in 1 of 999 draws its largest difference, 2.74, exceeds the observed 2.62.

The two placebos and the aligned placebo of Table 6 make different compromises. Table 6’s placebo deletes exactly as many pairs as the correction and matches both sale quarters, but stays in the original first-endpoint universe. The degree-matched placebo works in the stricter universe and nearly matches the pair count, but matches only the flagged deed’s quarter, not the quarter of the other end of each pair. For D4 the observed difference exceeds the maximum of every design except one draw of the degree-matched placebo; for D0 the margin is narrowest under Table 6’s placebo (3.69 against 3.41). The placebo margin is therefore sensitive to design. The findings do not contradict Table 6. Matching does not make placebo pairs exchangeable with flagged ones, so the share of draws above the observed value is a descriptive reference and not a calibrated p-value.

Unresolved classifications. Table 9 conditions on the classifier’s labels. Table 7 remains the analysis of the deeds the review leaves unresolved: its scenarios reassign deeds across the population at the review’s unresolved shares. Recomputed on the stricter universe, scenario A, which treats every third-party auction purchase as non-market, gives a largest difference of –3.45 in 2011Q3, against –3.37 in Table 7. It changes 4,162 pairs and moves the path by up to 1.50 log points relative to D0. An auction purchase by a third

party may be at a market price, so A overstates the non-market set by design. The bounds reported in Section 5, precision 0.893–0.993 and recall 0.135–0.523, refer to the review’s broad target: every transfer whose recorded amount is not a negotiated price, including family transfers. They are not bounds on the recall of 0.966 for foreclosure-type transfers, and they measure agreement with AI readings, not accuracy against a verified record.

What the appendix supports. Restricting both endpoints to 1–3 family houses changes the index path very little. The raw index lies above the drop series around 2010 and below it throughout 2011–12. For the published definition, the six-definition joint band excludes zero in five of the eight quarters of 2011–12. For the narrowest definition it excludes zero in none of them. Separately, the narrowest definition’s 2011–12 average has a bootstrap interval, not joint across definitions or quarters, that lies below zero. The size of the 2010Q3 peak depends on the definition, and its margin over a placebo depends on the placebo’s design. None of these checks establishes a change in the cumulative peak-to-trough decline or recovers a market price for any foreclosure conveyance.

Table 9: Raw minus drop for 1–3 family houses, pairs with two 1–3 family endpoints, log points

	D0 published	D4 narrowest
Pairs in universe	138,380	138,380
Flagged endpoint deeds	8,181	3,022
Pairs dropped	12,327	4,700
Difference in 2010Q3 (largest) [pointwise 95% CI]	+3.61 [+2.48, +4.98]	+2.62 [+1.72, +3.60]
Joint 95% band at 2010Q3, six definitions × 91 quarters	[+1.63, +5.59]	[+0.64, +4.60]
Quarters where the joint band excludes zero	10	2
Mean difference 2008–10 [95% CI]	+1.04 [+0.56, +1.56]	+1.05 [+0.85, +1.27]
Mean difference 2011–12 [95% CI]	–2.28 [–2.85, –1.74]	–1.08 [–1.30, –0.88]
Peak-to-trough difference [95% CI]	–0.28 [–3.04, +1.31]	+0.66 [–0.78, +1.07]
Deed placebo, borough × quarter: largest difference, median / 95th pct. / max	1.46 / 2.02 / 2.59	0.92 / 1.33 / 2.06
Deed placebo: mean pairs deleted; draws above observed	9,342; 0 of 999	3,555; 0 of 999
Degree-matched deed placebo: median / 95th pct. / max	—	1.24 / 1.74 / 2.74
Degree-matched: mean pairs deleted (range); draws above observed	—	4,664 (4,645–4,680); 1 of 999

Notes: Universe: the pairs of Table 4 whose first and second endpoint deeds are both classed as 1–3 family houses. Intervals come from 999 paired Poisson(1) parcel-bootstrap draws common to D0–D5 (seed 20260929), which differ from the 300 draws of Tables 5–7. Pointwise and window-mean intervals are 2.5th and 97.5th bootstrap percentiles for one definition and one quarter or window; neither is joint across definitions or quarters, and the pointwise interval at the selected quarter is not an interval for the maximum. The joint band is centered and unstudentized, with half-width 1.98 at every quarter; it holds jointly across D0–D5 and all quarters from 2003Q2 to 2025Q4, conditional on the labels. Peak is the maximum over 2005–08 and trough the minimum over 2008–13, as in Table 4. Deed placebos: 999 draws of ordinary endpoint deeds in this universe matched on borough × quarter counts (seed 20260930), plus, for the degree-matched version, on the number of eligible pairs each deed enters; no stratum runs short. “Draws above observed” is descriptive and not a p-value.

Declarations

Data availability. All inputs are public NYC Open Data resources and published index series retrieved from FRED; retrieval scripts, manifests and reconciliation counts are in the replication package (10.5281/zenodo.22925383). The companion papers' packages are deposited separately (Loschi 2026b, 2026a); no result in this paper is taken from them.

Funding. This research did not receive any specific grant from funding agencies in the public, commercial or not-for-profit sectors.

CRedit authorship contribution statement. Pablo Loschi: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Validation, Visualization, Writing – original draft, Writing – review and editing.

Declaration of competing interest. The author founded HolaCasaClara, a free comparator of Argentine mortgage loans. It has no connection to New York housing records or to the questions studied here. The author has no other competing interests to declare.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process. During the preparation of this work the author used OpenAI's ChatGPT and Codex, Anthropic's Claude and Google's Gemini to draft and revise text, develop and check code, and cross-check the replication. As part of the research, the document review in Section 5 used AI in place of human readers: ChatGPT read the deed images and Claude checked each reading against ACRIS open data and the protocol's rules; the buyer-name readings in the remark check are also AI assessments. After using these tools, the author reviewed and edited the content as needed, checked the scoring and the revision log, and takes full responsibility for the content of the publication.

References

- Anenberg, Elliot, and Edward Kung. 2014. "Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures." *American Economic Review* 104 (8): 2527–51. <https://doi.org/10.1257/aer.104.8.2527>.
- Bailey, Martin J., Richard F. Muth, and Hugh O. Nourse. 1963. "A Regression Method for Real Estate Price Index Construction." *Journal of the American Statistical Association* 58 (304): 933–42. <https://doi.org/10.1080/01621459.1963.10480679>.
- Ben-David, Itzhak. 2011. "Financial Constraints and Inflated Home Prices During the Real Estate Boom." *American Economic Journal: Applied Economics* 3 (3): 55–87. <https://doi.org/10.1257/app.3.3.55>.
- Campbell, John Y., Stefano Giglio, and Parag Pathak. 2011. "Forced Sales and House Prices." *American Economic Review* 101 (5): 2108–31. <https://doi.org/10.1257/aer.101.5.2108>.
- Case, Karl E., and Robert J. Shiller. 1987. "Prices of Single-Family Homes Since 1970: New Indexes for Four Cities." *New England Economic Review*, 45–56.
- City of New York. 2024. *New York City Administrative Code §§ 11-2101 to 11-2117 and Rules of the City of New York (19 RCNY § 23-03)*.
- Clauret, Terrence M., and Nasser Daneshvary. 2009. "Estimating the House Foreclosure Discount Corrected for Spatial Price Interdependence and Endogeneity of Marketing Time." *Real Estate Economics* 37 (1): 43–67. <https://doi.org/10.1111/j.1540-6229.2009.00234.x>.

- Conklin, James N., N. Edward Coulson, Moussa Diop, and Nuno Mota. 2023. "An Alternative Approach to Estimating Foreclosure and Short Sale Discounts." *Journal of Urban Economics* 134: 103546. <https://doi.org/10.1016/j.jue.2023.103546>.
- Cordell, Larry, Liang Geng, Laurie S. Goodman, and Lidan Yang. 2015. "The Cost of Foreclosure Delay." *Real Estate Economics* 43 (4): 916–56. <https://doi.org/10.1111/1540-6229.12107>.
- Doerner, William M., and Andrew Leventis. 2015. "Distressed Sales and the FHFA House Price Index." *Journal of Housing Research* 24 (2): 127–46. <https://doi.org/10.1080/10835547.2015.12092100>.
- Ellen, Ingrid Gould, Amy Ellen Schwartz, Ioan Voicu, and Michael H. Schill. 2007. "Does Federally Subsidized Rental Housing Depress Neighborhood Property Values?" *Journal of Policy Analysis and Management* 26 (2): 257–80. <https://doi.org/10.1002/pam.20247>.
- Federal Reserve Bank of St. Louis. 2026. *FRED Economic Data: Series NYXRNSA, ATNHPIUS35614Q, ATNHPIUS36005A, ATNHPIUS36047A and ATNHPIUS36081A*. <https://fred.stlouisfed.org/>.
- Fisher, Lynn M., Lauren Lambie-Hanson, and Paul Willen. 2015. "The Role of Proximity in Foreclosure Externalities: Evidence from Condominiums." *American Economic Journal: Economic Policy* 7 (1): 119–40. <https://doi.org/10.1257/pol.20130102>.
- Gerardi, Kristopher, Eric Rosenblatt, Paul S. Willen, and Vincent Yao. 2015. "Foreclosure Externalities: New Evidence." *Journal of Urban Economics* 87: 42–56. <https://doi.org/10.1016/j.jue.2015.02.004>.
- Harding, John P., Eric Rosenblatt, and Vincent W. Yao. 2009. "The Contagion Effect of Foreclosed Properties." *Journal of Urban Economics* 66 (3): 164–78. <https://doi.org/10.1016/j.jue.2009.07.003>.
- Kermani, Amir, and Francis Wong. 2026. "Racial Disparities in Housing Returns." *American Economic Review* 116 (1): 287–331. <https://doi.org/10.1257/aer.20240327>.
- Loschi, Pablo. 2026a. *The Shape of the Tax: Transaction Shifting at New York City's 2019 Transfer-Tax Notches and the Mortgage Recording Tax*. Zenodo. <https://doi.org/10.5281/zenodo.22925301>.
- Loschi, Pablo. 2026b. *Who Gets the House? Related-Party Transfers and the Cash–Mortgage Price Gap in New York City*. Zenodo. <https://doi.org/10.5281/zenodo.22421850>.
- New York City Department of Finance, Office of the City Register. n.d. *ACRIS Bandwidth Notice*. <https://a836-acris.nyc.gov/BandwidthPolicy/ACRIS-BW-POL.html>.
- New York State Department of Taxation and Finance. 2019. *Form TP-584.1: Real Estate Transfer Tax Return Supplemental Schedules, Schedule E and Instructions (Revision 7/19)*. https://www.tax.ny.gov/pdf/current_forms/property/tp584_1.pdf.
- Nowak, Adam D., and Patrick S. Smith. 2020. "Quality-Adjusted House Price Indexes." *American Economic Review: Insights* 2 (3): 339–56. <https://doi.org/10.1257/aeri.20190337>.
- Pennington-Cross, Anthony. 2006. "The Value of Foreclosed Property." *Journal of Real Estate Research* 28 (2): 193–214. <https://doi.org/10.1080/10835547.2006.12091177>.
- Schwartz, Amy Ellen, Scott Susin, and Ioan Voicu. 2003. "Has Falling Crime Driven New York City's Real Estate Boom?" *Journal of Housing Research* 14 (1): 101–35. <https://doi.org/10.1080/2167034x.2003.12461362>.
- S&P Dow Jones Indices. 2025. *S&P Cotality Case-Shiller Home Price Indices Methodology*. <https://www.spglobal.com/spdji/en/documents/methodologies/methodology-sp-cotality-cs-home-price-indices.pdf>.

U.S. Federal Housing Finance Agency. n.d. *FHFA House Price Index: All-Transactions Indexes*. <https://www.fhfa.gov/data/hpi>.