

Who Gets the House?

Related-Party Transfers and the Cash–Mortgage Price Gap in New York City

Pablo Loschi*

14 October 2026

Working paper, version 3.8

Abstract

Repeat sales can hold a property fixed while comparing transactions with different financing labels. Using New York City deed and mortgage records for 2016–2025, I find a house price contrast of 9.25 log points [7.17, 11.35] between purchases with and without a matched mortgage, after permit, institutional-seller, party-type and ZIP-year controls. The contrast varies with the cash buyer’s recorded identity. For company cash buyers it is 1.53 log points [-1.79, 5.44]; the interval includes zero and allows a positive gap. Excluding 358 pairs with a shared seller-buyer surname reduces the house contrast to 4.87 log points [2.69, 6.76]. Transfer reports more often show non-market signs among flagged deeds. The surname screen is a proxy and its effect does not identify a causal contribution of family transfers. Condominium contrasts are small under a common design. AI-assisted document checks identify missed mortgages. A fresh stratified probability sample of 40 endpoints finds purchase borrowing behind two unmatched-mortgage labels; 18 endpoints remain unresolved. Correcting the two supported labels changes the headline contrast from 9.25 to 9.23 log points. The readings do not establish classifier accuracy. Sensitivity scenarios preserve a positive gap under random recoding but permit reversals when errors depend on price growth. The contribution is evidence on the composition and measurement of a cash–mortgage contrast. The design identifies neither a causal financing effect nor the incidence of the mortgage recording tax.

JEL: G21, G51, R21, R31. *Keywords*: cash purchases; mortgage finance; related-party transfers; execution risk; repeat sales; New York City.

*Independent researcher, Berlin, Germany. Correspondence: loschi.pablo@gmail.com. ORCID: 0009-0004-9455-4713. This research received no external funding. The author founded HolaCasaClara, a free comparator of Argentine mortgage loans, which has no connection to New York housing records or this study. No other competing interests or outside rights to review the research before dissemination are reported.

1 Introduction

Buyers who pay cash often pay less. Across United States markets the mortgage–cash price difference is on the order of ten percent (Reher and Valkanov 2024), and Han and Hong (2024) link a smaller Los Angeles cash discount to the risk that a financed closing fails. The usual reading is that sellers accept less in exchange for certainty. The alternative is composition: cash purchases may differ in the properties, circumstances and parties involved, so that the average gap says more about who pays cash than about what certainty is worth. The distinction matters for policy. Proposals to tax cash purchases, or to offset the cash buyer’s advantage, take the gap as a measure of that advantage.

Hedonic studies put the cash discount at about 13 percent (Asabere et al. 1992) and find that it varies across periods once selection into cash purchase is modelled (Tidwell et al. 2018). Buyer and seller characteristics shift negotiated prices (Harding et al. 2003). Who buys also affects the price paid: investors drawn into the Los Angeles market during the boom paid more when buying and received less when selling than other investors (Bayer et al. 2021), and large buy-to-rent firms entered the single-family market after 2012 (Mills et al. 2019). Financing terms also move what buyers will pay: in a survey experiment, a lower required down payment raises stated willingness to pay, by far the most for renters (Fuster and Zafar 2021). Reher and Valkanov (2024) show that differences in sellers’ circumstances explain about half of the part of the national premium that a representative-seller model cannot, with belief distortions a candidate for the rest; sellers’ reference points and financial constraints also shape the prices they ask (Andersen et al. 2022). On the buyer side, administrative financing labels add a channel of their own: transfers between relatives are often recorded as unfinanced sales at reduced prices, so they enter the cash side of a financing contrast.

This paper measures the gap in New York City, where public deed and mortgage records can be linked to repeat sales and where financed buyers also bear a statutory mortgage recording tax of 1.80 to 1.925 percent of loan principal. An archived panel covers house and condominium sales in Manhattan, the Bronx, Brooklyn and Queens between 2016 and 2025. The financing contrast compares the appreciation of properties that switch from a cash to a financed purchase with those that switch the other way. It is estimated jointly with permit restrictions, institutional-seller exclusions, buyer and seller type controls and ZIP-year effects, and every change in specification is evaluated with a common parcel bootstrap.

For houses the contrast is 9.25 log points, with bootstrap interval [7.17, 11.35], and between 5.15 and 7.66 log points under estimators that limit the weight of extreme appreciation. Record-consistency restrictions, borough omissions and alternative financing windows leave it positive throughout. Condominiums behave differently: under a common borough-quarter

design their contrast is 0.77 log points $[-0.01, 1.56]$, and slightly negative under robust estimators.

The main result is heterogeneity in the measured contrast across cash-buyer categories. When the buyer at the unfinanced sale is a company, the contrast is 1.53 log points $[-1.79, 5.44]$. When it is an individual or another party, it is 13.77 $[10.96, 15.98]$. Deed party names supply a screen for one possible source of this composition. Sales on which the grantor and grantee share a surname, a marker of possible transfers between relatives, which need not be priced at arm's length, make up 5.7 percent of unfinanced sale endpoints and 2.3 percent of financed ones. Because such transfers sit on the cash side of both switching arms, they raise measured appreciation in cash-to-financed pairs and lower it in financed-to-cash pairs. If those transfers carry reduced prices, they can inflate the contrast while leaving the symmetry diagnostic of Section 4.2 unchanged. Excluding the 358 pairs with a same-surname deed changes the house contrast by -4.38 log points $[-5.63, -3.18]$, to 4.87 $[2.69, 6.76]$. Ignoring the 40 most common surnames to limit coincidental matches gives 6.71 $[4.69, 8.72]$. The flag is a proxy: it misses relatives with different surnames and can match unrelated people, and the change it produces is not an estimate of the share of the gap caused by family transfers. A reading of the state transfer reports filed with a random sample of deeds supports it: 40 percent of flagged deeds $[26, 55]$ declare a sale between relatives, a buyer who is also a seller or a non-standard deed, or report a zero price with no transfer tax, against 10 percent $[3, 30]$ of other cash deeds. Applied to both property types under a common design, the screen moves the house contrast from 11.42 to 7.50 log points and the condominium contrast from 0.77 to -0.46 $[-1.16, 0.14]$, so the divergence between property types survives. Also excluding deeds on which the grantor and grantee share a mailing address lowers the house contrast further to 3.76, but that additional change, -1.11 $[-2.76, 0.29]$, is not distinguishable from zero.

The divergence between property types and between cash buyers bears on the execution-risk reading. A seller indifference condition converts a contrast into the loss a seller would have to expect on a failed financed sale if the entire contrast were a premium for execution certainty. At an assumed 10 percent failure probability, the headline house contrast implies a loss of 87 percent of the price, and the contrast without same-surname transfers 45 percent, compared with about 7 percent for condominiums. These figures describe what execution risk would have to explain, not its share, because failure probabilities and continuation values are unobserved. Financing-specific associations with mortgage rates and inventory, which an execution-risk mechanism would predict, do not survive financing-specific time trends.

The paper makes three contributions. It applies one repeat-sales design to houses and condominiums in the same city and finds the contrast concentrated in houses. It shows

that the house contrast depends on the identity of the cash buyer, and that excluding same-surname deeds, which administrative financing labels inherit from the deed record, reduces the measured contrast by almost half. And it evaluates each threat (renovation, distress, party composition, local appreciation, record errors and extreme growth) by refitting the joint model with paired uncertainty. Han and Hong (2024) observe listings, agreement dates and relisting, which this paper cannot. Repeat sales remove fixed property attributes (Bailey et al. 1963) but not changes in condition (Nowak and Smith 2020) or transaction circumstances, and distressed sales, which trade at substantial discounts (Campbell et al. 2011) and drive differences in realized returns (Kermani and Wong 2026), are addressed through institutional-name exclusions and a document-level pilot.

The paper does not observe contracts, contingencies or failed closings. Its financing labels come from record linkage rather than verified loan purpose, and its surname flag is a proxy for a relationship, not a legal finding; Section 7 collects these limits. It does not estimate the incidence of the mortgage recording tax or responses to New York’s 2019 transfer-tax thresholds.

A companion paper (Loschi 2026a) shows that some deeds recorded as ordinary sales in the same registry carry amounts set by foreclosure transfer-tax rules rather than negotiation. Applying that paper’s classifier to the deeds here changes the house contrast by about half a log point (Section 4.5); that paper reports 99.3 percent precision against resolved AI-assisted document readings, with unresolved cases retained in sensitivity analyses. This is not an independently verified accuracy rate. A second companion paper (Loschi 2026b) measures the statutory mortgage recording tax and transaction shifting at the 2019 thresholds. It reports no repeat-sales estimates, and the contrast measured here is not an estimate of that tax’s incidence.

2 Financing institutions in New York City

2.1 The recording tax and ownership form

New York Tax Law Article 11 taxes mortgages of real property (New York State 2026; New York City Department of Finance 2026c). Sections 253 and 253-a, together with the city schedule, produce standard borrower rates of 1.80 percent below \$500,000 of mortgage principal and 1.925 percent at or above that amount for the covered residential categories. A qualifying lender generally pays the separate 0.25 percent special additional tax; exemptions and instrument-specific credits can change the amount due. Section 253(2) excludes the first \$10,000 of principal for qualifying one- and two-family residences, reducing that component

by \$30. These rules are distinct from transfer taxes on the sale price.¹

For illustration, a standard \$720,000 mortgage used to buy a \$900,000 condominium implies a borrower recording charge of \$13,860 before any applicable instrument-specific credits. A cash purchase records no purchase mortgage. A co-operative share loan is secured by personal property rather than a mortgage on an individually owned real-property unit, so it does not attract the same tax. This does not make a co-operative purchase exempt from all transfer taxes or other closing charges.

The distinction is statutory. Sellers' prices, financing terms and ownership choices can respond, so statutory borrower liability is not an estimate of the ultimate economic burden. The same boundary creates a measurement problem: searching real-property mortgages alone cannot identify co-operative share financing. ACRIS personal-property filings must be investigated separately.

3 Data and record linkage

3.1 Samples and the boundary of replication

The repeat-sales input contains 290,840 archived observations across five boroughs. Excluding Staten Island leaves 166,907 house and 82,488 condominium sales. The exclusion reflects the coverage of the ACRIS linkage used here; it is not a claim that Staten Island has the same financing measurement.

The panel-construction code combines NYC Department of Finance annualized sales with ACRIS real-property master and legal records. Condominium billing lots can differ from the unit lots on which deeds and mortgages are recorded. The archived code recovers a unique deed through borough, block, consideration and date before assigning the unit parcel. The four archived input files are byte-identical to repository commit `f404e48320a81e3bfe20127af2ed7c2d9fb5268e` (Loschi 2026c). All estimates are recomputed from that archived panel; its record matches are documented but not independently re-verified.

New public extracts add DOB permit dates, PLUTO geography, ACRIS deed parties and consideration, co-operative transactions and personal-property filings. The source agencies are cited explicitly (New York City Department of Finance 2026d,b,a; New York City Department of Buildings 2026; New York City Department of City Planning 2026; Freddie Mac 2026). Source identifiers, preserved extracts, software versions and the exact computational

¹Primary sources: New York Tax Law §§253, 253-a, and the NYC Department of Finance mortgage recording tax schedule. The analysis does not reconstruct individual mortgage tax returns.

boundary are documented in the replication README. The replication run starts from the archived panel and rebuilds every enrichment and analysis from the preserved public extracts; the raw-to-panel linkage is documented but not re-executed.

Linkage quality has several dimensions. Bailey et al. (2020) distinguish match rates and sample representativeness from false-link and missed-link errors. Their historical-data error rates cannot be transferred to ACRIS. Here, agreement with archived flags checks computation, while the document readings and unresolved-case sensitivities address evidence about the financing labels. These checks do not establish population accuracy.

3.2 Financing and repeat-sale construction

For houses and condominiums, financing is the archived indicator of a matched purchase-window mortgage. The strict, base and wide windows are respectively $[-15, +30]$, $[-15, +90]$ and $[-15, +180]$ days around the sale, using document dates with recording-date fallback in the panel-construction code. Narrow, base and wide labels are retained. The absence of a matched mortgage is referred to as cash for brevity, but can include unobserved borrowing, missed linkage or instruments outside the rule. Assignments and satisfactions must not be counted as new purchase mortgages.

Within each eligible house or condominium parcel, sales are sorted by date and adjacent observations at least 1,095 days apart are retained. This yields 8,208 house pairs and 6,584 condominium pairs. All sales fall between 2016 and 2025; in the headline house sample, first sales run from January 2016 to December 2022 and resales from March 2019 to December 2025. Adjacency is established within the archived, already restricted panel; it does not prove that no excluded underlying transfer occurred. Repeated unit sales in co-operatives require an additional unit identifier and follow the broader adjacency audit described below.

The permit restriction removes a house pair if a recorded permit issuance lies between its two sale dates, including endpoints. It leaves 6,316 pairs. Permits measure recorded work, not its cost, completion or quality. Unpermitted work can remain, while some permitted work has little bearing on value. The restriction is consequently a condition sensitivity, not a constant-quality guarantee.

3.3 Parties, consideration and geography

Each house endpoint is assigned a deed only when the closest candidate document date within 45 days is unique. Ties remain unresolved. A consideration sensitivity additionally requires a positive deed amount within the greater of one dollar or one percent of the DOF sale price. This checks agreement across records without establishing that consideration

equals a negotiated retail value.

Grantor and grantee names supply entity, trust/estate and unknown-party indicators. Explicit company forms take precedence over trust tokens. The narrower institutional rule identifies named financial institutions; a broader alternative also uses generic lending and servicing terms. A pair is excluded when either party is flagged at either date. These rules identify institution-associated transactions, not legally adjudicated foreclosure, deed-in-lieu or REO status. A four-case instrument pilot, reported in Appendix D, distinguishes a referee foreclosure transfer from a subsequent lender resale and confirms that ordinary DEED metadata can cover foreclosure instruments. The pilot does not validate the remaining name classifications. Natural-person residual categories likewise do not establish beneficial ownership. The same party records supply the same-surname flag of Section 5.

PLUTO version 26v2 supplies a fixed geographic crosswalk for 8,167 house pairs. After the preceding cumulative restrictions, geography reduces the sample from 6,033 to 6,006 pairs. The full specification spans 137 ZIP codes and 52 community districts. Current identifiers provide location controls rather than historical measures of building condition. Of 1,227 ZIP-year endpoint cells, 353 have fewer than five endpoints; fine geographic adjustment comes with limited support in some cells.

4 Measuring financing-associated price gaps

4.1 The joint estimator

For pair i with sale dates a and z , define $y_i = \log P_{iz} - \log P_{ia}$. Let Q_i contain indicators for cash-to-financed (cf) and financed-to-cash (fc) transitions. The nuisance design X_i contains changes between sale dates in borough-quarter indicators, with optional local-year and party-type terms. The joint regression is

$$y_i = X_i\gamma + Q_i\beta + \varepsilon_i, \quad \hat{\pi}_J = \frac{1}{2}(\hat{\beta}_{cf} - \hat{\beta}_{fc}). \quad (1)$$

Same-financing pairs help estimate the market controls and provide the switching coefficients' reference. All fixed effects enter in first differences; no interpretation is attached to individual coefficients in a redundant nuisance design. Sparse least squares projects onto that design, with numerical convergence checks.

A simpler statistic residualizes only the outcome and compares switching-group means.

With M_X the residual maker and

$$c_i = \frac{1}{2} \left(\frac{1\{cf\}_i}{N_{cf}} - \frac{1\{fc\}_i}{N_{fc}} \right), \quad (2)$$

it is $\hat{\pi}_R = c' M_X y$. In contrast, the joint coefficients satisfy $\hat{\beta} = (Q' M_X Q)^{-1} Q' M_X y$. Residualizing y does not also residualize the financing indicators. The two estimands can therefore respond differently when market controls change. Table 1 reports both, but only equation (1) is the financing coefficient in a jointly controlled regression, and all comparisons use it.

Neither estimator is automatically causal. Within-property differencing removes fixed attributes but leaves time-varying condition, credit access, bargaining circumstances and selection into financing. Permits, recorded party types and local appreciation address observed parts of those differences. They do not make financing exogenous.

4.2 What symmetry diagnoses

Suppose the two residual switching means can be written as

$$\mu_{cf} = \pi + u_{cf}, \quad \mu_{fc} = -\pi + u_{fc}. \quad (3)$$

Then their half-difference is $\pi + (u_{cf} - u_{fc})/2$, while their sum $A = u_{cf} + u_{fc}$. A small A restricts the sum of the omitted components, not their difference. Opposing condition or composition effects can cancel in A while biasing the half-difference.

This connects the permit and institution analyses. Permits address recorded changes to the property; party records address who transacts and under what observable institutional circumstances. An institution acquiring title without a new mortgage can enter a cash group even when the transfer arises from debt resolution. A small aggregate asymmetry cannot distinguish that composition from renovation. Each restriction must therefore be evaluated through a refitted model and paired uncertainty, rather than treated as validation of the other.

4.3 Cumulative evidence and neighbourhood appreciation

Table 1 gives one specification sequence. The institutional exclusion increases the permit-free gap, whereas buyer and seller type controls reduce it. This difference matters: the name-based exclusion does not explain the positive house contrast, and the party-control change is not a causal decomposition of that contrast.

The geography-linked baseline holds the sample fixed before adding local-year controls.

Table 1: Cumulative house robustness

Specification	Pairs	Residual	Joint	Joint 95% CI
Original sample	8,208	11.14	11.42	[9.60, 13.19]*
Permit-free	6,316	9.34	9.63	[7.80, 11.56]*
+ Named institutions excluded	6,033	10.48	10.82	[8.96, 12.67]
+ Buyer/seller type controls	6,033	7.23	8.50	[6.58, 10.42]
Same geography-linked sample	6,006	7.56	8.87	[6.81, 10.84]*
+ Community-district/year	6,006	7.47	9.50	[7.66, 11.33]
+ ZIP/year	6,006	6.39	9.25	[7.17, 11.35]*

All estimates and intervals are log points. Every row includes borough-quarter effects. The last two rows are alternative geographic specifications added to the same linked baseline. Institutional exclusions use grantor and grantee name rules at both dates, not verified distress. Stars denote percentile intervals from 999 common parcel-bootstrap refits; other intervals use joint parcel influence functions (CR0). Intervals are pointwise. The residual and joint columns are different estimands.

Community-district-year and ZIP-year effects are alternatives, each supplementing borough-quarter effects. The final joint estimate is 9.25 log points, with bootstrap interval [7.17, 11.35]. The corresponding residual statistic is smaller, but is not the coefficient of financing in the same jointly controlled regression.

Table 2 evaluates changes using their covariance. ZIP-year effects change the joint estimate by 0.39 log points, with interval [-1.02, 1.60]. The data therefore do not establish that local appreciation explains a positive portion of the joint house gap. They also do not rule out every economically meaningful local-composition contribution. Current geography, sparse cells and omitted time-varying characteristics limit that conclusion.

Table 2: Paired changes in the joint house contrast

Comparison	Change	95% CI
Party controls	-2.32	[-2.98, -1.65]
Community-district/year	0.63	[-0.12, 1.38]
ZIP/year	0.39	[-1.02, 1.60]
Original to final	-2.17	[-4.02, -0.36]

Log points. Differences use cross-specification covariance on original parcel clusters. ZIP/year and original-to-final intervals use the common full-refit bootstrap; the others use the joint analytic covariance. Party and geographic comparisons hold the relevant estimation sample fixed. Original-to-final also changes the sample.




Figure 1: The residual-mean statistic and joint coefficient across controls

Point estimates and pointwise parcel CR0 intervals. The figure uses one analytic interval method throughout; the principal table also reports common full-refit bootstrap intervals. Community-district/year and ZIP/year are alternative final controls.

All specifications’ observation influences are summed within the original parcel and stacked before computing covariance. Excluded parcels receive zero influence in that specification. This preserves the covariance induced by overlapping samples. The 999-draw bootstrap resamples original parcel clusters and refits all nuisance controls in four principal specifications using the same draws. It confirms the main joint-coefficient comparison. The high-dimensional residual statistic has a shifted resampling distribution, another reason not to interpret its attenuation as an explained component. Intervals are pointwise and do not incorporate unobserved classification error.

4.4 Outlier-robust estimates

The joint contrast is a mean, and repeat-sales growth has a long upper tail. In the headline sample, 15 percent of cash-to-financed pairs more than doubled in price, against 4 percent of pairs financed at both sales. Appendix Table 12 shows the consequence: removing the 60 parcels with the largest positive influence, one percent of the sample, lowers the estimate to 4.68 log points. Because that deletion selects on the fitted outcome, it has no valid sampling interpretation. Table 3 instead reports estimators fixed in advance that limit the weight of extreme growth without reference to the financing labels.

Table 3: Outlier-robust estimates of the financing contrast

Estimator	Houses, full controls (6,006 pairs)			Borough-quarter only	
	Contrast	95% CI	Change from LS	Houses	Condos
Least squares (reference)	9.25	[7.17, 11.35]		11.42	0.77
Huber M-estimator	6.73	[5.32, 8.76]	-2.52 [-3.31, -1.14]	8.83	-0.49
Trim growth at 1st/99th pct.	7.66	[5.77, 9.53]	-1.59 [-2.76, -0.45]	9.23	-0.45
Trim growth at 2.5th/97.5th pct.	5.15	[3.55, 6.78]	-4.10 [-5.64, -2.53]	6.91	-0.58

Log points. Columns 2–4 use the headline specification of Table 1 (ZIP-year and borough-quarter effects, buyer and seller type controls). Columns 5–6 use the common specification of Table 4. The Huber estimator ($k = 1.345$; MAD scale from least-squares residuals, updated once) refits all nuisance controls by iteratively reweighted least squares. Trimming removes pairs whose annualized log growth lies outside the stated percentiles of the estimation sample; cutoffs ignore financing labels. Intervals are percentile intervals from the same 999 parcel-bootstrap draws as Table 1, re-estimating the Huber scale and trimming cutoffs in every draw; the change column is the paired difference within draws. Tuning constants are conventional defaults and were not varied.

The Huber estimator, which down-weights 34 percent of cash-to-financed pairs, gives 6.73 log points [5.32, 8.76]. Trimming annualized growth at the 1st and 99th percentiles gives 7.66, and at the 2.5th and 97.5th percentiles 5.15. For houses, every robust estimate is below least squares and every interval excludes zero. The house contrast remains positive under these alternative estimators, but its magnitude is sensitive to extreme appreciation. Under the common borough-quarter specification, the Huber estimate is 8.83 [7.39, 10.28] for houses and -0.49 [-1.00, -0.03] for condominiums. The condominium contrast is much smaller and sensitive to the estimator, and the property-type divergence survives.

4.5 Property types and classification sensitivity

Table 4 puts houses and condominiums under common borough-quarter controls. The condominium estimate is 0.77 log points, with interval [-0.01, 1.56]. Even with common controls, differences in ownership form, properties and matching coverage prevent attributing the contrast to standardisation or tax status alone.

The house findings survive narrower and wider financing windows, stable-label restrictions, broader institutional exclusions and consideration-consistent deed matching (Table 5). Window disagreement is an internal consistency measure. Among house sales, strict versus base labels differ for 1,020 observations and wide versus base for 2,193; neither is a known error count. Similarly, passing a consideration match does not establish distress status. A stratified instrument-review queue accompanies the code; no false-positive rate is inferred from these consistency checks.

Table 4: House and condominium contrasts under common controls

Type	Pairs	N_{cf}	N_{fc}	Joint	95% CI
House	8,208	1669	1035	11.42	[9.72, 13.13]
Condo	6,584	1228	1659	0.77	[-0.01, 1.56]

Both rows use borough-quarter effects and adjacent eligible sales at least 1,095 days apart. All available eligible pairs are used; the cumulative house restrictions are not imposed on the condominium sample. Parcel-clustered CR0 intervals, in log points. Differences in coverage, ownership and financing measurement prevent a causal property-type interpretation.

Table 5: Classification and linkage sensitivity

Restriction	Pairs	Joint	95% CI
Financing window: strict	6,006	9.56	[7.76, 11.36]
Financing window: base	6,006	9.25	[7.44, 11.07]
Financing window: wide	6,006	8.60	[6.77, 10.42]
Stable labels at both dates	5,864	9.03	[7.18, 10.88]
Broader institution exclusion	5,918	9.31	[7.48, 11.13]
Unique deed / amount match	5,798	9.23	[7.44, 11.02]

All rows use the ZIP-year model, with borough-quarter and buyer/seller controls. Log points; parcel CR0 intervals. Stable labels agree between strict and wide windows at both dates. Amount consistency requires a unique nearest deed within 45 days and positive consideration within the larger of one dollar and one percent of the DOF price. These tests concern internal consistency, not classification accuracy.

Lender take-title deeds. The institutional rule removes pairs with a named financial institution at either sale, but a lender that takes title through foreclosure or a deed in lieu can also appear under a servicer, trust or REO-vehicle name. A companion paper (Loschi 2026a) shows that such deeds can record an amount set by transfer-tax rules rather than a negotiated price. Applying its party-name classifier, unchanged, to the deeds linked to the headline pairs flags 93 of the 6,006 pairs. Of their 96 flagged endpoints, 93 are unfinanced, so lenders taking title enter the analysis as cash buyers; 38 of the flagged pairs survive even the broader institution exclusion. Excluding the flagged pairs raises the house contrast to 9.78 log points [7.73, 11.84], a paired change of 0.52 [0.19, 0.97]. The classifier’s broader option gives 9.76, and combined with the same-surname screen of Section 5 the contrast is 5.35 [3.23, 7.35], against 4.87 for that screen alone. The direction is the one legal-rule amounts imply: a take-title amount above market value lowers measured appreciation in cash-to-financed pairs and raises it in financed-to-cash pairs, and both shrink the contrast. Against resolved AI-assisted readings, the companion paper reports precision of 99.3 percent and foreclosure-type recall of 96.6 percent; the exclusion remains a sensitivity rather than a correction. It does not change the paper’s conclusions.

5 Who pays cash?

5.1 Related-party transfers

A sale between relatives can be recorded at any price and seldom needs a purchase mortgage. For a financing contrast such transfers are worse than noise. A below-market unfinanced sale at the first date raises measured appreciation in a cash-to-financed pair; one at the second date lowers it in a financed-to-cash pair. Both movements enlarge the half-difference, and they offset in the sum A of equation (3), so the symmetry diagnostic cannot detect them. Campbell et al. (2011) likewise remove apparent intrafamily transfers from their Massachusetts deed data before measuring prices.

The preserved ACRIS party extract names the grantors and grantees of the deed linked to each endpoint. Linked deeds with party records cover 98 percent of first-sale and 99 percent of resale endpoints in the headline sample. A deed is flagged when a natural-person grantor and grantee share a surname, using names recorded in the ‘surname, given name’ form; company, trust and estate names never match. The flag marks 358 of the 6,006 pairs. Same-surname deeds are 2.5 times as common at unfinanced endpoints (5.7 percent) as at financed ones (2.3 percent). Their prices are consistent with below-market transfers. Cash-to-financed pairs with a flagged deed have mean log appreciation of 0.65, against 0.34 for

other cash-to-financed pairs; flagged financed-to-cash pairs fall by 0.27 on average, while other financed-to-cash pairs rise by 0.20.

Panel A of Table 6 re-estimates the headline specification without flagged pairs. The house contrast falls to 4.87 log points [2.69, 6.76], a paired change of -4.38 [-5.63, -3.18]. The Huber estimate falls to 4.02 [2.28, 5.73], and every estimator gives between 2.83 and 4.87 log points. The flag makes two kinds of error. Unrelated parties can share a common surname: ignoring the 40 most frequent surnames in the extract removes most such coincidences and gives 6.71 [4.69, 8.72], although it also misses genuine relatives with common names. Relatives with different surnames, including many spouses and in-laws, are never flagged, and endpoints without party records cannot be flagged. The contrast that remains is therefore likely to contain some related-party transfers.

Every sampled deed carries a New York State real property transfer report (form RP-5217NYC), on which the parties declare the price and any condition of the transfer. A random sample of 40 flagged and 20 unflagged cash-endpoint deeds (seed 20260924) was read from the ACRIS images with AI assistance, cross-checked against the transfer taxes paid, and documented deed by deed (Appendix F). 16 of the 40 flagged deeds, 40 percent [26, 55], declare a sale between relatives or former relatives (13 deeds), a buyer who is also a seller, or a deed that is not a warranty or bargain-and-sale deed, or report a zero price with no transfer tax; 9 more combine a shared address or an owner on both sides with no declared condition. Among the unflagged cash deeds the share is 10 percent [3, 30]. Declarations may omit relationships, while a non-standard deed alone does not establish non-market status. These are descriptive rates of disclosed signs, not validated lower bounds on non-market prevalence; the sample is too small to measure the flag’s error rates precisely.

5.2 Company and individual cash buyers

Panel B of Table 6 splits both switching arms by the buyer at the unfinanced sale. Companies, identified by an explicit corporate form, are the cash buyer in 794 of the 1,781 switching pairs. When a company pays cash, the contrast is 1.53 log points [-1.79, 5.44]. When another party does, it is 13.77 [10.96, 15.98], and 7.11 [4.48, 9.12] after excluding same-surname transfers. Even then the difference between the two groups is 5.56 log points [0.97, 9.07].

Company form does not establish institutional ownership: an LLC can represent an individual buyer. The estimated company-cash contrast is smaller, but its interval permits a positive gap as well as zero. The price advantage is concentrated among individual cash buyers, and excluding pairs with same-surname deeds removes almost half of the headline contrast. The remaining gap among individuals is not thereby identified as a premium for

Table 6: Related-party transfers and the identity of the cash buyer

<i>A. Financing contrast by estimator</i>						
Sample	Pairs	Least squares	Huber	Trim 1/99	Trim 2.5/97.5	
All pairs	6,006	9.25 [7.17, 11.35]	6.73 [5.32, 8.76]	7.66	5.15	
No same-surname deeds	5,648	4.87 [2.69, 6.76]	4.02 [2.28, 5.73]	4.03	2.83	
Same, strict flag	5,768	6.71 [4.69, 8.72]	4.94 [3.43, 6.79]	5.13	3.06	
<i>B. Least-squares contrast by the buyer at the cash sale</i>						
Sample	Pairs	Company buyer	Other buyer	Other minus company		
All pairs	6,006	1.53 [-1.79, 5.44]	13.77 [10.96, 15.98]	12.24	[7.50, 16.06]	
No same-surname deeds	5,648	1.55 [-1.75, 5.49]	7.11 [4.48, 9.12]	5.56	[0.97, 9.07]	

Log points; 95% percentile intervals from the paper’s 999 common parcel-bootstrap draws (seed 20260909), with every nuisance control, Huber scale and trimming cutoff re-estimated in each draw. Headline specification of Table 1. A same-surname transfer is a deed at either endpoint on which a natural-person grantor and grantee share a surname; company, trust and estate names never match. The strict flag ignores the 40 most frequent surnames in the party extract. Panel B splits both switching arms by whether the grantee at the unfinanced sale has an explicit company form; always-financed and always-unfinanced pairs are pooled.

certainty: individual cash purchases can also include estate sales, sales to tenants and other circumstances that names do not reveal.

5.3 Condominiums and a shared-address screen

Deeds and parties were also retrieved for condominium endpoints, with every indexed request checked against the source’s record count. A unique deed is linked for 13,092 of 13,168 condominium endpoints; 31 ambiguous matches are left unlinked. The same surname rule flags 189 of 6,584 condominium pairs. Panel A of Table 7 applies the screen to both property types under the common borough-quarter design of Table 4. The house contrast falls from 11.42 to 7.50 log points [6.00, 9.19]; the condominium contrast moves from 0.77 [-0.02, 1.55] to -0.46 [-1.16, 0.14], a paired change of -1.23 [-1.71, -0.82]. Ignoring the 40 most common condominium surnames gives -0.26 [-0.97, 0.34]. Screening both property types therefore leaves the divergence between them intact.

A second proxy uses the mailing addresses recorded for each party. Panel B of Table 7 additionally excludes 1,595 house pairs with a deed on which the grantor and grantee share a normalized street address, city and ZIP code. The contrast falls to 3.76 log points [1.16, 6.05], but the additional change of -1.11 [-2.76, 0.29] is not distinguishable from zero, and restricting the match to natural persons gives 4.41 [2.14, 6.75]. Shared addresses are about as common at unfinanced endpoints (16.7 percent) as at financed ones (17.2 percent) and

include common mailing agents and co-located parties, so this screen is a broad sensitivity rather than evidence of kinship.

Table 7: The same-surname screen across property types, and a shared-address sensitivity

<i>A. Common borough-quarter design (Table 4)</i>					
	Pairs	All pairs	Flagged	Unflagged pairs	Change
Houses	8,208	11.42 [9.60, 13.19]	449	7.50 [6.00, 9.19]	-3.93 [-4.81, -2.94]
Condominiums	6,584	0.77 [-0.02, 1.55]	189	-0.46 [-1.16, 0.14]	-1.23 [-1.71, -0.82]
<i>B. Headline house specification: excluding pairs with</i>					
	Pairs	Contrast	Change vs. surname screen		
Same-surname deeds	5,648	4.87 [2.69, 6.76]			
or shared address	4,053	3.76 [1.16, 6.05]	-1.11	[-2.76, 0.29]	
or shared address, persons	4,458	4.41 [2.14, 6.75]	-0.46	[-1.43, 0.71]	

Log points; 95% percentile intervals from the paper’s 999 common parcel-bootstrap draws (seed 20260909), refitting all controls in each draw. Panel A: least squares under borough-quarter effects only, as in Table 4; the same-surname flag of Table 6 is applied to deeds linked to each endpoint. Condominium deeds are linked for 13,092 of 13,168 endpoints (31 ambiguous matches left unlinked). Panel B: headline specification of Table 1. A shared address is a grantor and grantee on the same deed with the same normalized mailing street, city and ZIP code; it can reflect a common mailing agent or co-location and is not evidence of kinship. The natural-person variant counts only addresses of parties with personal names.

6 Interpreting the financing gap

6.1 What the magnitudes imply

Consider a stylised seller comparing a certain cash price P_c with a financed offer P_f . Let q be the probability that the financed transaction fails and let dP_c be the loss, relative to the cash alternative, conditional on failure. If continuation following failure is worth $(1 - d)P_c$, indifference implies

$$P_c = (1 - q)P_f + q(1 - d)P_c, \quad \frac{P_f - P_c}{P_c} = \frac{qd}{1 - q}. \quad (4)$$

For an observed log contrast π , treating that entire contrast as the model premium requires

$$d^*(q; \pi) = (e^\pi - 1) \frac{1 - q}{q}. \quad (5)$$

This inversion makes the required loss explicit while leaving q as an assumption. Mortgage application denial is not the failure probability of an accepted purchase contract, so national

denial rates are not substituted for q here.

Table 8: Relisting-loss scenarios implied by the observed contrasts

Failure probability q	Houses: $100d^*$	Condos: $100d^*$
0.1	87.2	7.0
0.2	38.8	3.1
0.3	22.6	1.8

Scenario calculations from $d^* = (\exp(\pi) - 1)(1 - q)/q$. Houses use the full geographic joint estimate; condos use the borough-quarter joint estimate. They have different covariate sets and populations. Values of q are illustrative assumptions, not estimated contract-failure probabilities. No common relisting technology is imposed as an empirical fact.

At any common failure probability, the house contrast implies a loss on failure more than ten times as large as the condominium contrast if the entire gap is attributed to execution risk. At an assumed $q = 0.1$, it would require a continuation loss of 87 percent of the price, or 63 percent at the Huber estimate, compared with about 7 percent for condominiums (and zero or negative under robust condominium estimates). Without same-surname pairs the house requirement falls to 45 percent, or 37 percent at the Huber estimate, which is still several times the condominium figure; the screened condominium contrast requires no loss at all. A seller whose financed buyer fails still owns the house and can relist it. While a loss of a few percent aligns with ordinary carrying and relisting costs, a loss exceeding half the property’s value illustrates the extreme demands that an execution-risk explanation places on the house contrast. These calculations illustrate model requirements under stated assumptions rather than estimating the actual mechanism share, because contract failure rates and continuation values are unobserved.

Two qualifications apply. First, equation (5) evaluates the counterfactual in which the entire observed contrast represents a causal financing premium, which the design does not establish. The calculation evaluates what execution risk would have to explain under that assumption, rather than bounding or estimating the true closing-risk contribution. Second, failure probabilities could differ between the two property types. To reconcile the least-squares contrasts with a common loss, the failure odds would have to be about twelve times higher for houses than for condominiums. The mechanism in Han and Hong (2024), where market-specific relisting prospects interact with closing risk, is the natural candidate for such a difference and would require contract-level data to test.

6.2 Comparison with vendor-filtered estimates

Reher and Valkanov (2024) estimate a mortgage-cash premium of about 11 percent across United States markets from 1980 to 2017 after dropping transactions that the data vendor flags as intra-family transfers or foreclosures. They put the premium that rational sellers facing mortgage frictions would require at about 3 percent. New York City’s free deed records carry no such flag. Here the surname screen plays that role, more crudely: dropping pairs in which seller and buyer share a surname lowers the house contrast from 9.25 to 4.87 log points, which lies between their rational benchmark and their headline, and the company-cash contrast, 1.53 log points, is below both. Part of a contrast estimated from public records can therefore reflect transfers that vendor data would remove before estimation. The comparison does not show that the remaining gap equals an execution-risk premium, and the samples, periods and estimators differ.

6.3 What the co-operative comparison can establish

Co-operative share financing is a possible institutional comparison, but the recorded loan measure differs from the house and condominium measure. Timing-unique filing assignments leave 1,765 repeat pairs with an interval including zero; an instrument pilot nevertheless finds a filing that pledges a different unit from its assigned sale. That sale contributes no repeat pair, so rejecting its link leaves the estimate unchanged. A second filing matches the sale unit but lacks independent confirmation of purchase-loan purpose. Appendix C reports the full assignment and adjacency audit. The co-operative comparison is therefore a measurement sensitivity, not a control group that identifies a mortgage-tax effect.

6.4 Credit conditions at both sale dates

If credit conditions shape financing risk, attaching a rate only to the resale misses conditions at the initial transaction. The full house design therefore adds separate terms $F_a(r_a - 4)$ and $F_z(r_z - 4)$, where r is the latest available weekly Freddie Mac 30-year PMMS rate, in percent, on or before each sale. Quarterly effects do not absorb all within-quarter weekly-rate variation. Separate uninteracted rates at both endpoints are therefore included explicitly. The exercise studies financing-specific rate associations.

The slopes are small and negative at the first date and positive at the second date (Table 9). Adding separate financing-specific linear calendar trends reverses both signs; the two rate interactions are not jointly distinguishable from zero in the calendar-dependence sensitivity (Appendix Table 11). Rate movements, time trends and financing composition

are difficult to disentangle in these data. The rate evidence consequently does not identify the execution-risk mechanism. The PMMS methodology change on 17 November 2022 is an additional measurement qualification. The augmented model’s reference contrasts at 4 percent and the trend origin must not be reported as an unconditional average financing gap.

Table 9: Credit conditions at both sale dates

Financing $\times$ rate	Without trends	SE	With trends	SE
First sale	-0.43	1.95	0.46	1.93
Second sale	2.98	0.62	-1.24	1.03

Log points per one percentage point of PMMS rates, centered on 4 percent; 6,006 pairs. Both models include separate uninteracted weekly rates at both endpoints, the full geographic controls and both financing-switch indicators. The second adds separate endpoint financing-specific linear calendar trends. Standard errors cluster by parcel (CR0). Appendix Table 11 reports calendar-dependence sensitivity. The headline contrast is estimated in the model without rate interactions.

A bounded extension attaches prior-month borough inventory to both sale dates using StreetEasy’s Single Family series (StreetEasy 2026). All 6,006 house pairs are covered. The resale-date financing interaction reverses sign when financing-specific trends are added, and the joint calendar-quarter test then does not reject zero endpoint interactions. Appendix E reports both specifications. This sensitivity does not isolate a liquidity or execution-risk mechanism.

7 Conclusion

In New York City, financed purchases of houses are associated with prices 9.25 log points above cash purchases of the same house, and between 5.15 and 7.66 log points under estimators that limit the weight of extreme appreciation. The contrast is smaller for company cash buyers: 1.53 log points [-1.79, 5.44]. That interval does not establish equal prices. The larger contrast is associated with the other recorded buyer categories. Excluding pairs with a deed on which the seller and buyer share a surname reduces the house contrast by almost half, to 4.87 log points [2.69, 6.76]. Screened the same way, condominium contrasts are close to zero, while houses retain a clear positive contrast. Under a simple seller indifference model with an assumed 10 percent failure probability, even the reduced house contrast would require a loss of 45 percent of the price on a failed sale to be a pure premium for execution certainty, against about 7 percent for condominiums.

These conclusions rest on measured, not verified, transaction attributes. Financing labels

come from matching purchase-window mortgages and are not confirmed loan purposes; the absence of a match is not proof of a cash purchase. A shared surname or mailing address is a proxy for a relationship, not proof of one: it misses relatives with different names and can match unrelated people. Transfer reports in a small random sample of deed images, read with AI assistance, show non-market signs in 40 percent of flagged deeds and 10 percent of other cash deeds; these declarations are self-reported and the sample is small. Permits capture recorded renovation, not all changes in condition. Institutional-name rules mix foreclosure transfers with lender resales and cannot establish distress; excluding lender take-title deeds flagged by the companion classifier raises the house contrast by 0.52 log points. Co-operative financing assignments rely on timing and can attach a filing to the wrong unit. The panel’s record linkage is archived and documented but was not rebuilt from raw records for this version. A document review of 53 sampled pairs, read by AI, confirms 54 of 56 financed labels and finds a purchase mortgage behind 4 of 50 cash labels; dropping the 4 affected pairs leaves the house contrast at 9.40 log points (Appendix G). Under assumed random errors in either direction, at rates chosen for the sensitivity exercise, the average contrast would stay between 5.82 and 8.96 log points and the surname screen would still remove about half of it. Errors related to price growth matter more: the contrast would reach zero only if the recoded labels were concentrated on unfavourable pairs by 0.70 standard deviations on average (0.45 without same-surname pairs), and the review is too small to rule that out. All reported intervals condition on these measurements.

For policy, the results caution against reading an observed cash–mortgage price difference as the value of closing certainty or as the advantage of investors who pay cash. In these data, little of the gap belongs to company buyers, and a measure that includes family transfers is a poor basis for any tax or subsidy rate on the payment method. In May 2026, New York legislators considered a buyer-paid tax of 1 percent on New York City home purchases of \$1 million or more made without a mortgage, presented as parity with the mortgage recording tax; a week later it was reported as likely to be dropped from the state budget (Bloomberg News 2026a,b). The results here do not evaluate that proposal, which would apply only above \$1 million, but they caution against justifying such a tax by an observed cash discount. Research that uses administrative financing labels should screen related-party transfers before interpreting a cash discount. Separating execution risk from selection requires data on accepted offers, contingencies, failed closings and relisting, or an exogenous source of financing variation. The statutory mortgage recording tax is a distinct and directly measurable cost to financed buyers; its size and the design of New York’s transfer-tax thresholds are examined in Loschi (2026b).

Declarations

Data availability. All inputs are public records of New York City agencies and Freddie Mac. Code, archived inputs, preserved public extracts and complete outputs are deposited at doi:10.5281/zenodo.22421850 (concept DOI, resolving to the current version). A master script regenerates every table and figure from those files; the replication README documents the one stage not re-executed, the original raw-to-panel linkage. The companion papers’ replication packages are deposited separately (Loschi 2026b,a); no result in this paper is taken from them.

Funding and interests. Sole author; independent researcher. There was no external funding. The author founded HolaCasaClara, a free comparator of Argentine mortgage loans, which has no connection to New York housing records or this study. No other competing interests or outside rights to review the research before dissemination are reported.

Prior dissemination. Earlier versions circulated as a working paper (SSRN 7421138) and with the replication deposit. The paper is not under review elsewhere.

Generative AI. Generative AI tools (OpenAI’s ChatGPT and Codex, Anthropic’s Claude and Google’s Gemini) assisted with drafting, literature checking, code development and replication review. The author is responsible for the manuscript, interpretations and citations. The instrument review in Appendix G used AI in place of human readers: ChatGPT read the deeds and mortgages, and Claude then checked each reading against ACRIS open data; the author checked the scoring. The fresh 40-endpoint sample in Appendix G.1 was read with Codex, with a second AI instrument reading for ten cases; it has no independent human validation. The research uses administrative public records; no human subjects were involved.

A Joint uncertainty and numerical implementation

Write $r = M_X y$ and $m = M_X c$. The derivative of the residual contrast with respect to an observation weight, evaluated at unit weights, is

$$\text{IF}_{R,i} = m_i r_i - \frac{1\{cf\}_i \bar{r}_{cf}}{2N_{cf}} + \frac{1\{fc\}_i \bar{r}_{fc}}{2N_{fc}}. \quad (6)$$

The first term accounts for refitting the nuisance projection; the remaining terms account for changes in group denominators. For the joint estimator, let $R = M_X Q$ and $e = M_X y - R\hat{\beta}$. Its coefficient derivative is

$$\text{IF}_{\beta,i} = (R'R)^{-1} R'_i e_i. \quad (7)$$

Apply the half-difference contrast, sum within parcel and stack across specifications. The covariance is the cross-product of these stacked cluster influences (CR0). Paired-difference variances use the corresponding covariance terms, rather than treating the two estimates as independent.

A numerical check perturbs five actual parcel weights in the full ZIP-year model, refits the nuisance parameters and verifies both analytic derivatives against finite differences with absolute error below 10^{-7} . The common bootstrap draws multinomial weights over the 8,127 original house parcels, using seed 20260909, and refits the original, permit-free, linked-party-control and ZIP-year specifications. All 999 draws are retained. Redundant nuisance columns are projected out without interpreting an arbitrary individual fixed-effect solution.

No multiple-testing correction is applied; all intervals are pointwise. These covariance calculations condition on the observed linkage and classifications.

B Replication map and classification definitions

The master script `code/reproduce.py` parses the preserved rate extracts, rebuilds enriched pairs, estimates the house models and bootstrap, audits co-operative assignments, estimates the outlier-robust models, performs numerical checks and regenerates manuscript exhibits. The replication README contains exact setup instructions, package versions and a file-level exhibit map. Archived original inputs remain distinguishable from newly retrieved extracts and derived outputs.

Table 1, Table 2 and Figure 1 derive from the repeat-sale results and joint bootstrap outputs. Table 4 recomputes both property types from the archived panel. Table 3 is produced by `code/robust_estimators.py` from the same parcel-bootstrap draws. The co-operative, credit and classification tables map to their named analysis scripts. Table 8 is

computed directly from the displayed formula and identified point estimates.

The loan windows in the house/condominium input are document-based purchase-window proxies; the co-operative windows concern personal-property recording dates. These do not have identical coverage or timing interpretation. The party classifications are reproducible regular-expression categories, not verified persons, beneficial owners or distressed sales. Unknown party records remain explicit rather than being silently classified as individuals. The release retains audit fields and review queues so subsequent adjudication can be versioned without overwriting the underlying rules.

C Co-operative assignment and adjacency audit

C.1 Why the co-operative comparison requires an audit

A building-level UCC filing is not inherently a particular unit buyer’s purchase loan. Multiple sales and filings can occur near each other, some sales lack recovered unit identifiers, and filings can concern other credit events. The co-operative extract therefore includes all observed prices and incomplete-unit records when counting possible competitors for a filing, rather than using only the estimation sample.

The broadest rule, the expanded count rule, works as follows: for a sale at date d , count eligible sales within $d \pm 45$ days and filings within $[d - 60, d + 135]$. A zero filing count is coded cash; a filing count at least as large as the sale count is coded financed; the remainder is ambiguous. The expansion is consequential. Of 60,496 purchases labelled financed, 4,385 have no filing in the base $[-15, +90]$ -day window. Counting all observed competing sales also changes assignments.

The alternative rules require exactly one candidate filing and exactly one observed sale eligible to claim that filing. The narrow, base and wide windows are $[-15, +30]$, $[-15, +90]$ and $[-15, +180]$ days. Documents linked to multiple parcels are excluded. These conditions prevent reuse of a supposedly unique filing, but cannot identify its debtor, collateral unit or purchase purpose. Zero candidates still means no matched record, not proven cash.

Adjacency is determined using all observed known-unit sales before dropping endpoints with ambiguous financing or out-of-band prices. Forming pairs within the estimation sample alone would bridge an intervening observed unit sale in 271 cases, so Table 10 forms adjacency on all known-unit sales in every row. Base timing uniqueness leaves 1,765 pairs in 1,082 buildings. Its interval includes zero, as do the narrow and wide alternatives. This is compatible with a small co-operative contrast within selected, differently measured samples; it is not evidence that the mortgage recording tax causes the house gap. An instrument

pilot makes the remaining linkage problem concrete. Among two timing-unique candidates inspected, one filing names the recorded sale unit; the other pledges a different unit in the same building. The latter sale contributes no repeat pair, so rejecting its filing link does not change Table 10. The matching case still lacks independent confirmation of purchase-loan purpose. These observations demonstrate a failure mode, not a population error rate, and preclude treating timing uniqueness as validated financing.

Table 10: Co-operative assignment sensitivity

Assignment rule	Pairs	N_{cf}	N_{fc}	Joint	95% CI
Expanded count rule	5,839	247	602	-1.18	[-3.11, 0.75]
All competing sales	5,261	229	576	-0.89	[-2.77, 1.00]
Unique timing: narrow	2,840	441	537	-0.41	[-2.00, 1.18]
Unique timing: base	1,765	219	252	-1.42	[-3.70, 0.86]
Unique timing: wide	1,265	108	143	0.09	[-2.99, 3.17]

Log points and building-clustered CR0 intervals. Every row forms adjacency using all observed known-unit sales before dropping ineligible endpoints and ambiguous financing. Timing uniqueness does not establish debtor identity or purchase purpose. No matched filing is a cash proxy, not proof of an unlevered purchase. Estimates use the September 2026 co-operative extract.

D Targeted instrument pilot

Four queued cases were inspected through the public ACRIS image viewer on 10 September 2026: one broad-only house flag, one specific-institution house flag, and two timing-unique co-operative assignments. Saved page images, document identifiers, review notes and source hashes accompany the replication package. These are AI-assisted document readings. The earlier 90-house and 90-co-operative queues remain incomplete and are preserved unchanged; their sampling weights must not be applied to this selectively continued four-case pilot to estimate error rates. The focused queue of Appendix G was instead completed with AI-assisted readings and an open-data cross-check; it is not the independent two-human-reader review originally planned.

House deed 2023051600851001 recites the lender’s acquisition under referee deed 2019020500740001. The earlier operative instrument describes a mortgage foreclosure and conveyance to the lender; the later transaction is consistent with a lender resale. The broader name rule flags the sale whereas the narrower rule does not. House deed 2019093000695001 is explicitly a referee’s deed in foreclosure and is already caught by the

narrower institutional-buyer rule. Both source cover records use DEED metadata. These findings distinguish mechanisms within institutional flags without measuring property condition or the effect of distress on prices.

For co-operative filing 2025020500312001, the sale unit and collateral designation both identify 7AB, but purchase-loan purpose remains unverified. Filing 2020020400882001 pledges unit 4E, whereas its timing-assigned sale identifies 1CD. The preserved sale extract contains only one observation for that latter unit; the false link therefore affects no repeat pair. It is rejected as evidence for that sale’s financing, without relabelling the sale as cash. A timing-only rule cannot resolve this unit mismatch even when it has no observed competing sale.

E Historical market conditions and shared calendar shocks

The inventory extension uses one selected indicator: the log base two of StreetEasy’s prior-month borough sales inventory in its Single Family category. It is centered on each borough’s unweighted mean log inventory over December 2015–November 2025, so a unit represents a doubling relative to that reference. The downloaded September 2026 vintage covers both endpoints of all 6,006 pairs. Its listing coverage need not represent the administrative one-to-three-family house sample. Inventory reflects both supply and demand, including platform participation; it is not a direct measure of closing risk. Historical revisions and unknown publication dates mean that a prior-month observation is not necessarily information publicly available on the sale date.

Both endpoint inventory levels enter uninteracted, alongside separate financing-inventory interactions, the full house controls, and endpoint financing-by-borough controls with Brooklyn omitted. The latter prevent fixed geographic differences in financing associations from being attributed to inventory movements. Separate financing-specific linear calendar trends are then added as a sensitivity. This plan was recorded after checking source coverage and before fitting these associations; it was not externally preregistered. Both fitted specifications are reported, without choosing a preferred indicator or trend specification from significance.

Parcel-only uncertainty is supplied in the replication output. The table additionally allows score dependence whenever two repeat pairs share a parcel or either endpoint calendar period, regardless of whether the shared period belongs to their first or second sales. Inclusion–exclusion counts overlapping memberships once. The calendar-quarter calculation uses 40 nodes; the calendar-year sensitivity uses only ten. Negative eigenvalues in the esti-

Table 11: Historical market conditions: endpoint and trend sensitivity

Model	First	95% CI	Second	95% CI	p_Q	p_Y
Credit	-0.43	[-5.91, 5.05]	2.98	[2.02, 3.94]	< 0.001	< 0.001
Credit + trends	0.46	[-4.63, 5.54]	-1.24	[-3.18, 0.70]	0.418	0.330
Inventory	-3.13	[-7.17, 0.91]	-13.38	[-22.54, -4.21]	0.005	0.003
Inventory + trends	-4.51	[-9.79, 0.77]	3.21	[-6.65, 13.08]	0.195	0.244

All 6,006 pairs; slopes in log points. Credit unit: one percentage point of the weekly rate. Inventory unit: a doubling of prior-month borough Single Family inventory relative to its borough reference. Inventory models also include separate endpoint financing-by-borough controls (Brooklyn reference). Both indicators enter uninteracted at both dates. Intervals use the union of parcel and either-endpoint calendar-quarter score dependence and a t_{39} reference. Columns p_Q and p_Y test both interaction slopes jointly, using calendar-quarter and calendar-year dependence respectively, with approximate $F_{2,39}$ and $F_{2,9}$ references. Negative covariance eigenvalues are set to zero: one is corrected in each quarter model with trends and in each year model. The unadjusted eigenvalues are recorded in the replication output. These are exploratory, pointwise diagnostics, not exact small-cluster inference.

mated covariance are replaced by zero, with their magnitudes retained in the output. The resulting t and F references are approximate diagnostics, not exact small-cluster tests. These calculations do not cover arbitrary dependence across non-overlapping periods, and intervals are pointwise without multiplicity adjustment.

With trends, the inventory interactions are not jointly distinguishable from zero under either calendar aggregation. The analogous credit exercise reaches the same qualitative conclusion after common weekly rates are included. This does not demonstrate that market conditions are irrelevant; it shows that these associations are sensitive to their time specification. The headline conditional house contrast is estimated in the unchanged main model and is not replaced by the reference contrast from either interaction model.

F Deed-image review of the same-surname flag

The review sample was drawn by a seeded script before any deed was read: 40 endpoint deeds at random among flagged pairs and 20 among unflagged cash endpoints of the headline house sample (seed 20260924). All 684 pages of the 60 recorded documents were retrieved from the public ACRIS image service. Each contains a cover page, the deed, and a New York State real property transfer report (RP-5217NYC). Item 12 of that report gives the full sale price and item 14 lists conditions of the transfer, including (A) a sale between relatives or former relatives, (C) a buyer who is also a seller and (E) a deed type other than warranty or bargain and sale. Prices and conditions were read from the report images; cover pages supplied the parties' addresses and the city and state transfer taxes, whose rates imply the

taxed consideration. Where the report price and the taxed consideration disagree, both are recorded and the case is classed as unclear rather than resolved. Readings were made with AI assistance and documented deed by deed. A complete independent human review has not been established. The readings, image crops and scripts are in `audit/deed_image_review`.

A deed is classed as showing a non-market sign when it declares condition A, C or E, or reports a zero price with no transfer tax. The classification is conservative. Conditions are self-declared: a transfer from a father to a son of the same name, for instance, declares none. Among flagged deeds, 9 further cases list a buyer among the sellers or give the same address for both parties without declaring a condition. One flagged deed shares only a common surname among seven parties and is dropped by the strict flag, and one unflagged deed declares a sale between relatives with different surnames. The review therefore supports the flag as a marker of non-market transfers without measuring its sensitivity or specificity.

G Focused validation of the house contrast

A bounded validation pass checks the exact archived rows generating the 6,006 analysis pairs. There are 5,957 parcels and 11,963 distinct sale endpoints. All pair price changes and financing flags reconstruct from those rows, and the three financing windows are nested. The linear coefficient weights reproduce the joint house contrast to numerical precision. These checks establish computational consistency, not the correctness of the archived source linkage.

Ten pairs involve an endpoint with more than one house record on the same parcel and date. Three financed sale endpoints have no positive recorded mortgage amount. Neither finding proves a false sale or financing classification, but both warrant explicit scrutiny. Omitting these cases and applying the existing safeguards jointly leaves 5,579 pairs and a contrast of 9.02 log points. The archived house panel does not retain the mortgage instrument identifiers needed for independent purchase-loan verification.

Table 12 reports the complete refitted sequence.

Influence is calculated from the fitted joint model and aggregated by parcel. Every deletion refits all nuisance controls. Removing the single largest absolute-influence parcel leaves 9.07 log points; the 60-parcel absolute-influence deletion leaves 8.51. This argues against a single-record explanation. It does not establish insensitivity to adversarial sample selection: removing the 60 largest positive influences lowers the estimate to 4.68. Section 4.4 addresses the same concern with robust estimators using conventional tuning constants that have valid intervals. Omission of each borough and each of the three largest ZIP groups is reported, without selecting the most favorable comparison.

The replication output also reports adversarial perturbations to pair log-price growth,

Table 12: Focused house validation: influence, geography and source consistency

Restriction	Pairs	Contrast	Change
Main reference	6,006	9.25	+0.00
Omit top 1 absolute influences	6,005	9.07	-0.18
Omit top 5 absolute influences	6,001	9.68	+0.42
Omit top 10 absolute influences	5,996	9.79	+0.54
Omit top 30 absolute influences	5,976	9.41	+0.16
Omit top 60 absolute influences	5,946	8.51	-0.74
Omit 60 largest positive influences	5,946	4.68	-4.57
Omit Manhattan	5,962	9.26	+0.01
Omit Bronx	4,962	9.42	+0.17
Omit Brooklyn	4,052	10.39	+1.14
Omit Queens	3,042	7.91	-1.34
Omit ZIP 11234	5,801	9.11	-0.14
Omit ZIP 10465	5,844	9.01	-0.25
Omit ZIP 11357	5,875	9.25	-0.00
Joint record-consistency safeguards	5,585	8.98	-0.27
Omit ambiguous source dates	5,996	9.29	+0.04
Omit missing positive loan amounts	6,003	9.25	-0.00
Joint safeguards and source checks	5,579	9.02	-0.23

Log points. Every row refits the joint switch model and all nuisance controls on its retained sample. Parcels, not single observations, are removed in the influence exercises. Absolute- and positive-influence exclusions are selected using the fitted outcome; they are diagnostics, not alternative preferred samples, and ordinary confidence intervals would not account for that selection. Geographic and source restrictions also change the estimation population. The three ZIPs are selected by pair count. Combined safeguards require stable financing windows, unique amount-consistent deeds at both endpoints and no broad institutional-name flag. Source checks additionally remove ambiguous parcel/date records and financed endpoints without positive recorded loan amounts. These are internal consistency restrictions, not validated classifications. The complete model outputs and nominal parcel intervals are retained in the replication files.

holding the labels and design fixed. It selects parcels with the largest absolute coefficient-weight sums and moves their outcomes against the weight signs. These exact linear calculations quantify sensitivity to outcome contamination under stated assumptions. They are not financing-misclassification bounds or estimates of unobserved renovation effects.

A price-blinded review queue samples 20 pairs at random from each financing-switch arm, separately from 13 targeted cases: the pairs on the ten most influential parcels, five in each arm, and the three pairs with a financed endpoint lacking a positive mortgage amount. The queue therefore contains 53 pairs and 106 sale endpoints. The original protocol called for two human readers; this version uses AI instead. ChatGPT read the deed and every candidate mortgage for each of the 106 endpoints, recording the borrower, the collateral, the purpose and the pages cited. Claude then checked each reading against ACRIS open data, covering parties, amounts and every mortgage recorded on the parcel from 15 days before to 90 days after the sale, and returned the inconsistent readings for a second look. The readings were locked before the analyst key was opened. Claude did not read the images, so the two steps are not independent readings, and the rates below measure agreement with the readings rather than accuracy against a verified record. Because the original financing labels were built from the same public mortgage records, the open-data check tests consistency rather than providing a fully independent source.

The documents confirm a purchase mortgage for 54 of the 56 endpoints labelled financed; of the other two, one shows a new loan on a transfer in which a buyer is also a seller and the other two land loans to the buying company. These are conflicts, not evidence of cash. Among the 50 endpoints labelled cash, 4 have a same-day purchase mortgage in the documents (8 percent of the reviewed cash-labelled endpoints). The pooled count includes targeted cases and has no population-representative error interval. These are documented label errors. Two arise from a date error in the city’s index: in one pair the mortgage, and in the other the deed, is indexed a year away from its true date, so the matching window missed the loan. The other two are loans recorded against two adjoining lots. A missing loan cannot prove a cash purchase, and no reading establishes one, so no cash label can be confirmed and pair-level error bounds stay wide: in the two probability-sampled switch arms, between 6.2 and 100 percent of pairs carry a label error, with all 3 observed errors in the financed-to-cash arm.

Dropping the 4 pairs with a documented error and refitting the headline model gives 9.40 log points (analytic parcel-clustered interval [7.59, 11.21]), against 9.25 [7.44, 11.07] for the published sample on the same interval. Treating as unresolved one loan recorded 25 days after the transfer and a two-lot development loan leaves 3 documented errors and 9.39 log points. Whether hidden purchase mortgages raise or lower the contrast depends on

which pairs carry them: correcting the documented errors raises it slightly, while the random recoding below lowers it slightly. The readings, the open-data check and the scoring code are in `audit/cross_ai_review/` and `code/cross_ai_review.py`.

The review cannot say how many unreviewed cash labels hide a purchase mortgage. Table 13 therefore asks how far the contrast would move if similar errors existed elsewhere. It starts from the corrected sample (9.40 log points; 4.89 without same-surname pairs), holds the 49 remaining reviewed pairs at their reviewed labels, and recodes at random a share of the other cash labels to financed: of the 712 unreviewed financed-to-cash pairs and of the 1,019 unreviewed cash-to-financed pairs. The shares are scenarios informed by the detected discrepancies, not measured error rates: 2 or 3 of the 20 sampled financed-to-cash pairs, depending on whether the two-lot development loan counts, and none of the 20 cash-to-financed pairs, up to the Wilson upper limits of 36 and 16 percent. Each draw fits the same recoded data with and without same-surname pairs. Across 200 draws per scenario, the mean contrast stays between 9.13 and 9.34 log points, and between 4.69 and 4.87 without same-surname pairs. In the most pessimistic scenario the 2.5th percentiles are 7.92 and 3.56. The surname screen removes 48–49 percent of the contrast on average in every scenario.

These ranges are distributions of point estimates under assumed random, one-direction recoding, not confidence intervals, and they omit sampling and validation uncertainty. Table 14 relaxes the three restrictions of this test.

Both directions and the whole sample. Panel A recodes unreviewed labels anywhere in the sample and in either direction: a share of the 2,613 cash-labelled endpoints to financed, at assumed rates of 8 percent or 18.8 percent, motivated by the pooled reviewed count and its nominal Wilson upper limit, and a share of the 9,293 financed-labelled endpoints to cash, at 3.6 percent (the two unresolved financed readings counted as errors) or 6.4 percent (the nominal Wilson upper limit for no detected error in 56). These pooled endpoints include targeted cases, so the rates are scenario settings rather than population estimates or design-based confidence limits. Cash–cash and financed–financed pairs can now become switching pairs, and the arms are rebuilt. Random recoding of this kind pulls the contrast toward zero, to between 5.82 and 8.96 log points on average, because it mixes pairs of typical growth into both arms. It does not change the sign, and the surname screen still removes 47–51 percent of the contrast.

Documented failure modes. Panel B follows the two mechanisms the review found. An index-date error or a late recording shows up as a mortgage in the wide window but not the base window; recoding every such unreviewed cash endpoint to financed (90 endpoints) gives 8.75 log points. Drawing the hidden mortgages only from company buyers, whose development loans can cover several lots, gives 9.26 to 9.39.

Errors related to growth. Whether a recoded label raises or lowers the contrast depends on the pair’s growth net of the model’s time and location effects. A hidden mortgage in a financed-to-cash pair with weak growth, for example, removes that pair from its arm and lowers the contrast. Panel C draws the 544 recoded labels with a probability that rises with their effect against the contrast, and reports that relationship as the mean standardized effect of the labels chosen, which is zero when errors are unrelated to growth. The contrast falls steadily and reaches zero when that mean is 0.70 standard deviations; without same-surname pairs it reaches zero at 0.45. Choosing the labels with the largest effect in either direction bounds the estimate between -15.70 and 25.49 log points. These bounds are wide because a few hundred labels can be placed on the pairs that matter most. They describe what the data cannot exclude, not a likely error process. The documented failure modes in Panel B show no such pattern, but the review is too small to rule it out.

Table 15 applies the documented corrections and the screens together on the paper’s common bootstrap. Removing the documented errors changes the contrast by $+0.15$ $[-0.03, 0.36]$; the surname screen then removes about half of it; the broad lender-name rule changes it by -0.01 ; and excluding legal-rule deeds changes it by $+0.31$ $[0.11, 0.54]$. With all of them the contrast is 5.19 log points $[2.98, 7.18]$ on 5,520 pairs, a change of -4.06 $[-5.36, -2.76]$ from the published estimate.

Together these checks show what the labels can support. The contrast and the share removed by the surname screen hold under random errors in both directions and under the error patterns the review documents. They would not hold if errors were concentrated on the pairs whose growth works against the contrast, which the review can neither detect nor rule out.

A targeted document follow-up. A later pass read the deeds, transfer reports and candidate mortgages of seven influential endpoints flagged by a tabular trace of all 11,963 sale endpoints. It found two problems. Three sale prices cover several parcels: ten lots in one case and two in the others, one of them an industrial lot in a sale that the transfer report marks as commercial. No allocation to the house appears in the pages read, so these pairs do not compare the prices of one house. Four cash endpoints have a mortgage whose index date is one or two years earlier than the dates written in the mortgage and the note, which match the buyers’ deed. The matching window missed these loans, as it missed the two index-date errors that the review above found. Purchase financing is inferred from the same-day deed, mortgage and collateral; no use-of-proceeds clause was read. The pass was AI-assisted, unblinded and only partly confirmed by the author, so it supports the handling of these cases but measures no error rate. The cross-reference for case P02 has

Table 13: Conditional stress test: random recoding of unreviewed cash labels

Share recoded, financed-to-cash arm	Share recoded, cash-to-financed arm	All pairs	Without same- surname pairs	Share of contrast removed by screen
Corrected baseline, no recoding		9.40	4.89	48%
0.10 (2 of 20)	0.00	9.33 [8.81, 9.82]	4.87 [4.41, 5.31]	48% [45, 51]
0.10 (2 of 20)	0.08	9.26 [8.57, 9.93]	4.78 [4.11, 5.44]	48% [45, 52]
0.10 (2 of 20)	0.16	9.20 [8.47, 9.86]	4.75 [4.01, 5.35]	48% [45, 54]
0.15 (3 of 20)	0.00	9.34 [8.66, 9.97]	4.87 [4.28, 5.40]	48% [44, 52]
0.15 (3 of 20)	0.08	9.24 [8.53, 10.00]	4.79 [4.10, 5.51]	48% [44, 52]
0.15 (3 of 20)	0.16	9.22 [8.47, 10.06]	4.78 [4.06, 5.60]	48% [43, 53]
0.36 (upper limit)	0.00	9.17 [8.04, 10.30]	4.78 [3.72, 5.90]	48% [41, 56]
0.36 (upper limit)	0.08	9.13 [8.04, 10.13]	4.69 [3.67, 5.77]	49% [41, 56]
0.36 (upper limit)	0.16	9.13 [7.92, 10.40]	4.76 [3.56, 5.90]	48% [40, 57]

Headline joint model, log points. Baseline: the 6,002 pairs left after removing the four pairs with a documented label error (Appendix G); 5,645 without same-surname pairs. Each draw recodes, at random, the stated share of the 712 unreviewed financed-to-cash pairs and of the 1,019 unreviewed cash-to-financed pairs from cash to financed at their cash endpoint; the 49 reviewed pairs keep their labels. Both columns are fitted on the same recoded data. Entries are means over 200 draws, with the 2.5th and 97.5th percentiles of the point estimates in brackets; these are scenario ranges, not confidence intervals. Shares are assumed scenarios informed by the review’s detected discrepancies (2 of 20 under the strict coding, 3 of 20 under the lenient one; 0.36 and 0.16 are Wilson upper limits), not measured error rates. `code/financing_stress_test.py`.

Table 14: Financing-label errors in both directions, across the whole sample and related to price growth

Scenario	Labels recoded	Arms (c-to-f / f-to-c)	All pairs	Without same- surname pairs	Removed by screen
Corrected baseline, no recoding	0	1,044 / 733	9.40	4.89	48%
<i>A. Random, both directions, all pairs</i>					
Hidden 8.0%, false financed 0.0%	209	995 / 711	8.96 [8.29, 9.67]	4.72 [3.99, 5.43]	47%
Hidden 8.0%, false financed 3.6%	544	1,101 / 821	7.34 [6.31, 8.17]	3.80 [2.84, 4.63]	48%
Hidden 18.8%, false financed 0.0%	491	924 / 671	8.24 [7.16, 9.22]	4.32 [3.18, 5.18]	48%
Hidden 18.8%, false financed 6.4%	1,086	1,115 / 884	5.82 [4.64, 6.82]	2.98 [1.86, 3.82]	49%
<i>B. Documented failure modes</i>					
Late or early recording, all	90	1,044 / 715	8.75	4.27	51%
Company buyers only, hidden 8.0%	209	1,003 / 719	9.26 [8.60, 9.92]	4.83 [4.09, 5.51]	48%
Company buyers only, hidden 18.8%	491	935 / 694	9.39 [8.42, 10.31]	4.93 [3.95, 5.98]	48%
<i>C. Related to growth, hidden 8.0%, false financed 3.6%</i>					
Mean standardized effect 0.31	544	–	4.12	1.15	–
Mean standardized effect 0.56	544	–	1.52	–0.88	–
Mean standardized effect 0.81	544	–	–1.18	–2.86	–
Mean standardized effect 1.19	544	–	–5.17	–5.79	–
Worst case against the contrast	544	1,054 / 809	–15.70	–14.28	9%
Worst case in its favour	544	1,217 / 860	25.49	20.75	19%

Headline joint model, log points, on the corrected sample of Table 13; the 49 reviewed pairs keep their labels. Hidden: share of the 2,613 unreviewed cash-labelled endpoints recoded to financed; false financed: share of the 9,293 unreviewed financed-labelled endpoints recoded to cash; at most one endpoint per pair. Arms are rebuilt after recoding. Panels A and B: means over 200 draws with the 2.5th and 97.5th percentiles of the point estimates; late or early recording recodes every cash endpoint with a mortgage in the wide but not the base window. Panel C draws the recoded labels with probability rising in their effect against the contrast; the first column gives the mean standardized effect of the labels recoded (0 when unrelated to growth), 96 draws each. Worst cases pick the labels with the largest effect. Scenario ranges, not confidence intervals. `code/misclassification_checks.py`.

Table 15: Documented corrections and screens applied together

Sample	Pairs	Contrast [95% CI]	Change from published	Change from row above
Published headline sample	6,006	9.25 [7.17, 11.35]	–	–
Less the four documented label errors	6,002	9.40 [7.23, 11.41]	+0.15 [–0.03, 0.36]	–
and same-surname pairs	5,645	4.89 [2.73, 6.79]	–4.36 [–5.62, –3.17]	–4.51 [–5.74, –3.30]
and the broad lender-name rule	5,558	4.88 [2.74, 6.83]	–4.37 [–5.65, –3.10]	–0.01 [–0.39, 0.45]
and legal-rule deeds	5,520	5.19 [2.98, 7.18]	–4.06 [–5.36, –2.76]	+0.31 [0.11, 0.54]

Headline joint model, log points. Each row removes the stated pairs from the row above. Intervals are 2.5th–97.5th percentiles over the paper’s common parcel bootstrap (999 draws, seed 20260909), so every row and difference is paired with the published estimate. Flags in the published sample: 358 same-surname pairs, 88 pairs matching the broad lender-name rule and 93 pairs with a legal-rule deed (Section 5, Appendix G). `code/misclassification_checks.py`.

since been reconciled by reopening the referenced instrument. Document 2022042100383001 is a power of attorney, not a competing deed. The endpoint deed is 2022042100383002, and the operative mortgage and note are dated 7 April 2022; the mortgage cover instead indexes 7 April 2020 and cross-references the power of attorney. This resolves the document-identity question without identifying loan use from an explicit proceeds clause. The date-rule recordings remain sensitivity scenarios, not newly verified purchase-financing labels. Its findings are in `audit/targeted_followup`.

The seven cases were chosen for their influence, so correcting only them could favour one direction. Table 16 therefore also applies each problem as a rule to every pair. Dropping the 80 pairs with a multi-parcel endpoint deed gives 9.82 log points, and recoding the 45 cash endpoints that meet the index-date rule gives 9.36. Together they give 9.91 [8.09, 11.73], close to the 9.85 obtained by correcting the seven cases alone. Applying both corrections raises the estimated contrast slightly. The index-date rule is narrower than the late-or-early recording scenario of Table 14, which recodes every cash endpoint with a mortgage in the wide window and gives 8.75: it requires the loan to be recorded near the sale, to name the buyer and to cover only the purchased parcel. The headline estimate is unchanged; `code/targeted_followup.py` reproduces the table.

Table 16: Targeted document follow-up and the same rules applied to every pair

Handling	Pairs	Contrast	Analytic 95% interval
Published sample	6,006	9.25	[7.44, 11.07]
<i>A. Seven targeted cases</i>			
Drop three pairs with a bundle price	6,003	9.66	[7.88, 11.44]
Recode four index-date conflicts as financed	6,006	9.44	[7.62, 11.25]
Both	6,003	9.85	[8.07, 11.63]
Both, with the earlier review's documented handling	6,001	9.99	[8.21, 11.76]
<i>B. The same rules applied to every pair</i>			
Drop pairs with a multi-parcel endpoint deed (120 endpoints)	5,926	9.82	[8.03, 11.61]
Recode cash endpoints under the index-date rule (45 endpoints)	6,006	9.36	[7.52, 11.21]
Both	5,926	9.91	[8.09, 11.73]

Log points. Intervals are analytic and parcel-clustered, conditional on each scenario; they differ from the bootstrap intervals elsewhere in the paper and omit classification uncertainty. Panel A: the earlier review's handling drops two pairs and recodes one endpoint, leaving one case unresolved. Panel B, index-date rule: a single-parcel mortgage against the parcel, recorded from 15 days before to 90 days after the sale, with an index date outside that window and a recorded mortgagor matching the deed's grantee.

G.1 A fresh probability sample of financing labels

A separate draw on 9 October 2026 sampled previously unreviewed endpoints from the headline population. After excluding 111 endpoints covered by earlier reviews, 11,852 unique endpoints remained: 2,598 without a matched mortgage and 9,254 with one. The seeded draw selected 20 without replacement from each stratum, then hid the recorded labels in an analyst key. The queue, key and input hashes were frozen. Readings were locked before the key was opened. This sample complements the earlier targeted checks and does not represent all New York City sales.

A frozen parcel-and-date search found candidate instruments for 22 endpoints. AI readings examined the operative deeds and mortgage instruments for all 22. Ten received a second AI instrument reading, including the three cases initially supported only by recording covers, seller financing and date, parcel or transfer-type questions. These are not independent human readings. One instrument expressly says purchase money; the other purchase connections are inferred from the acquisition, new loan, buyer-borrower match, collateral and timing. The other 18 endpoints had no candidate in the frozen search and remain unresolved. A failed search does not establish an unfinanced acquisition.

After unlocking, all 20 recorded-financed endpoints had supporting mortgage evidence. Two of the 20 recorded-unfinanced endpoints also had it. One mortgage has a January 2022 index date but a January 2023 operative date matching the acquisition. The other secures the purchased house together with a separate parking parcel. Both proposed recordings had received a second instrument reading before unlocking. The date conflict and the bundled price are retained as separate issues.

Allowing unresolved cases to agree or disagree gives a design-weighted sample bound of 2.19–21.92 percent for financing-label discrepancy in the previously unreviewed population. Exact finite-population hypergeometric inversion within each stratum, with a Bonferroni allocation across four one-sided tails, gives a conservative 95 percent confidence envelope of 0.19–37.26 percent. These calculations are conditional on the AI readings being correct. They include sampling and unresolved-case uncertainty, but cannot quantify correlated reading mistakes. The broad envelope does not support a precise population error rate.

Table 17 corrects only the two sampled labels, retaining the original sample and model. The contrast changes from 9.25 to 9.23 log points. Excluding the house-and-parking bundle gives 9.26; also excluding the sampled lender resale and referee auction gives 9.28. The intervals are conditional analytic parcel-clustered intervals. No unsampled label is extrapolated from this small sample, and the main estimate remains unchanged. The readings, frozen design, comparison and executable analysis are in `audit/probability_sample_20261009/`; its `second_read/analyze_probability_sample.py` reproduces the calculation.

Table 17: Fresh probability sample: source coverage and limited corrections

<i>A. Twenty endpoints from each recorded-financing stratum</i>			
Recorded label	Population	Supported mortgage	Unresolved
Without matched mortgage	2,598	2	18
With matched mortgage	9,254	20	0
<i>B. Sensitivity in the original headline model</i>			
Handling	Pairs	Contrast	Analytic 95% interval
Original headline sample	6,006	9.25	[7.44, 11.07]
Correct two sampled financing labels	6,006	9.23	[7.42, 11.05]
Also exclude house-and-parking bundle	6,005	9.26	[7.44, 11.07]
Also exclude two special-transfer cases	6,003	9.28	[7.47, 11.10]

Panel A covers previously unreviewed unique endpoints in the 6,006-pair headline population; sample sizes are 20 in each stratum. Supported means an AI reading of acquisition and mortgage evidence. Unresolved cases are finite-search negatives, not confirmed cash purchases. Panel B reports log points and analytic parcel-clustered intervals, conditional on the stated recodings and exclusions; these differ from the main bootstrap and omit classification uncertainty. The two special transfers are a lender resale and a referee auction conveyance. Their exclusion is an additional scenario, not a conclusion that their prices are unusable. Archived data and headline results remain unchanged.

References

- Steffen Andersen, Cristian Badarinza, Lu Liu, Julie Marx, and Tarun Ramadorai. Reference dependence in the housing market. *American Economic Review*, 112(10):3398–3440, 2022. doi: 10.1257/aer.20191766.
- Paul K. Asabere, Forrest E. Huffman, and Seyed Mehdian. The price effects of cash versus mortgage transactions. *Real Estate Economics*, 20(1):141–154, 1992. doi: 10.1111/1540-6229.00576.
- Martha J. Bailey, Connor Cole, Morgan Henderson, and Catherine Massey. How well do automated linking methods perform? lessons from US historical data. *Journal of Economic Literature*, 58(4):997–1044, 2020. doi: 10.1257/jel.20191526.
- Martin J. Bailey, Richard F. Muth, and Hugh O. Nourse. A regression method for real estate price index construction. *Journal of the American Statistical Association*, 58(304):933–942, 1963. doi: 10.1080/01621459.1963.10480679.
- Patrick Bayer, Kyle Mangum, and James W. Roberts. Speculative fever: Investor contagion in the housing bubble. *American Economic Review*, 111(2):609–651, 2021. doi: 10.1257/aer.20171611.
- Bloomberg News. NY plans new tax on NYC homes over \$1 million bought with cash. Bloomberg, 14 May 2026, 2026a. URL <https://www.bloomberg.com/news/articles/2026-05-14/ny-plans-new-tax-on-city-homes-over-1-million-bought-with-cash>.
- Bloomberg News. NYC cash home purchases tax is likely to be dropped from New York state budget. Bloomberg, 21 May 2026, 2026b. URL <https://www.bloomberg.com/news/articles/2026-05-21/new-york-s-tax-on-nyc-cash-home-purchases-is-close-to-collapse>.
- John Y. Campbell, Stefano Giglio, and Parag Pathak. Forced sales and house prices. *American Economic Review*, 101(5):2108–2131, 2011. doi: 10.1257/aer.101.5.2108.
- Freddie Mac. Primary mortgage market survey: Historical archives, 2026. URL https://www.freddiemac.com/pmms/pmms_archives. Weekly 30-year mortgage rates, 2015–2025; 574 observations parsed from preserved annual archive extracts.
- Andreas Fuster and Basit Zafar. The sensitivity of housing demand to financing conditions: Evidence from a survey. *American Economic Journal: Economic Policy*, 13(1):231–265, 2021. doi: 10.1257/pol.20150337.
- Lu Han and Seung-Hyun Hong. Cash is king? understanding financing risk in housing markets. *Review of Finance*, 28(6):2083–2118, 2024. doi: 10.1093/rof/rfae025.
- John P. Harding, Stuart S. Rosenthal, and C. F. Sirmans. Estimating bargaining power in the market for existing homes. *Review of Economics and Statistics*, 85(1):178–188, 2003. doi: 10.1162/003465303762687794.
- Amir Kermani and Francis Wong. Racial disparities in housing returns. *American Economic Review*, 116(1):287–331, 2026. doi: 10.1257/aer.20240327.

- Pablo Loschi. When a deed is not a market sale: Foreclosure transfers, statutory consideration, and repeat-sales house-price measurement. Working paper and replication package, Zenodo, 2026a. URL <https://doi.org/10.5281/zenodo.22925383>. Concept DOI resolving to the latest version.
- Pablo Loschi. The shape of the tax: Mortgage recording liability and transaction shifting at New York City’s 2019 transfer-tax thresholds. Working paper, Zenodo, 2026b. URL <https://doi.org/10.5281/zenodo.22925301>. Concept DOI resolving to the latest version. Replication package: doi:10.5281/zenodo.22880953.
- Pablo Loschi. Who gets the house? replication code and archived inputs, 2026c. URL <https://github.com/pgold30/who-gets-the-house>. GitHub snapshot, commit f404e48320a81e3bfe20127af2ed7c2d9fb5268e; accessed September 10, 2026.
- James Mills, Raven Molloy, and Rebecca Zarutskie. Large-scale buy-to-rent investors in the single-family housing market: The emergence of a new asset class. *Real Estate Economics*, 47 (2):399–430, 2019. doi: 10.1111/1540-6229.12189.
- New York City Department of Buildings. DOB permit issuance, 2026. URL <https://data.cityofnewyork.us/Housing-Development/DOB-Permit-Issuance/ipu4-2q9a>. NYC Open Data: ipu4-2q9a; preserved original permit snapshot.
- New York City Department of City Planning. PLUTO, version 26v2, 2026. URL <https://www.nyc.gov/site/planning/data-maps/open-data/dwn-pluto-mappluto.page>. NYC Open Data: 64uk-42ks; targeted parcel geography extract, September 2026.
- New York City Department of Finance. ACRIS personal property master and legals, 2026a. URL <https://www.nyc.gov/site/finance/property/acris.page>. NYC Open Data: sv7x-dduq and uqqa-hym2; September 2026 extract.
- New York City Department of Finance. ACRIS real property master, legals, and parties, 2026b. URL <https://www.nyc.gov/site/finance/property/acris.page>. NYC Open Data: bnx9-e6tj, 8h5j-fqxa, 636b-3b5g; archived panel and September 2026 targeted extracts.
- New York City Department of Finance. Mortgage recording tax, 2026c. URL <https://www.nyc.gov/site/finance/property/property-mortgage-recording-tax-mrt.page>.
- New York City Department of Finance. Annualized calendar sales, 2026d. URL <https://data.cityofnewyork.us/City-Government/NYC-Citywide-Annualized-Calendar-Sales-Update/w2pb-icbu>. NYC Open Data, w2pb-icbu; archived 2016–2025 price sample and September 2026 co-operative extract.
- New York State. New York Tax Law §§ 253, 253-a, 261, 1402, 1402-a, 1402-b, 1404; New York City Administrative Code §§ 11-2101 to 11-2117. Consolidated Laws of New York, Chapter 60, Articles 11 and 31; Real Property Transfer Tax, 2026. §1402-b added by L.2019, c.59, pt. D, effective 1 July 2019, with conveyances under binding contracts entered into on or before 1 April 2019 excluded.
- Adam D. Nowak and Patrick S. Smith. Quality-adjusted house price indexes. *American Economic Review: Insights*, 2(3):339–356, 2020. doi: 10.1257/aeri.20190337.

Michael Reher and Rossen Valkanov. The mortgage-cash premium puzzle. *Journal of Finance*, 79(5):3149–3201, 2024. doi: 10.1111/jofi.13373.

StreetEasy. Streeteasy data dashboard: Historical total sales inventory, single family. <https://streeteasy.com/blog/data-dashboard/>, 2026. Borough-month series, downloaded 22 September 2026; retrospective source vintage.

Alan Tidwell, Andres Jauregui, Vivek Sah, and Andrew Narwold. Cash and distressed house sales price discounts: Dual sample selection spatial interdependence approaches. *Journal of Real Estate Finance and Economics*, 56(1):101–139, 2018. doi: 10.1007/s11146-016-9591-y.