Who Gets the House?

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

Pablo Loschi*

29 October 2026

Working paper, version 3.12

Abstract

Cash buyers often pay less for homes, and in 2026 New York legislators considered taxing cash purchases of New York City homes. I ask how much of the gap between cash and financed prices reflects who transacts. I link sales in four New York City boroughs to deed parties and mortgage records from 2016 to 2025. In 6,006 house pairs, a repeat-sales model gives a conditional financed-minus-cash contrast of 9.25 log points (95 percent interval [7.17, 11.35]). Excluding the 358 pairs in which seller and buyer share a surname removes 47 percent of it, leaving 4.87 [2.69, 6.76], while deleting as many matched unflagged pairs produces no similar decline. When the cash buyer is a company the contrast is smaller and imprecise, at 1.53 [-1.79, 5.44], and condominium contrasts are much smaller under common controls. Cash here means no matched mortgage, and a shared surname is only a proxy for a relationship. The results show how transaction screens and the identity of the cash buyer change a financing contrast measured from public records. They do not identify the causal value of paying cash or the share of the gap due to family transfers.

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 about 11 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: 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, and it was left out of the state's enacted budget (Bloomberg News 2026a,b; Office of the New York City Comptroller 2026). Arguments of this kind take the observed cash discount as a measure of the cash buyer's advantage. Both of the estimates above come from vendor data that drop flagged family transfers and foreclosures before estimation. Public records carry no such flag, and I measure how much that composition changes the gap.

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). In Orange County, Florida, buyers who were later foreclosed paid up to about 3 percent more near the 2007 peak, with second-home buyers and small investors paying more while investor-companies and owner-occupiers did not (Turnbull and van der Vlist 2022). 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, and sellers' reference points and financial constraints also shape the prices they ask (Andersen et al. 2022). Administrative financing labels add a channel of their own on the buyer side. Transfers between relatives are often recorded as unfinanced sales at reduced prices, and they then enter the cash side of any financing contrast built from those labels.

I measure 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. My archived panel covers house and condominium sales in Manhattan, the Bronx, Brooklyn and Queens between 2016 and 2025, and the financing contrast I use compares the appreciation of properties that switch from a cash to a financed purchase with those that switch the other way. I estimate it jointly with permit restrictions, institutional-seller exclusions, buyer and seller type controls and ZIP-year effects, and evaluate every change in specification 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. It stays positive under record-consistency restrictions, borough omissions and alternative financing windows. Condominiums behave differently: under a common borough-quarter design their contrast is 0.83 log points (analytic interval [0.06, 1.61]), and it turns slightly negative under robust estimators.

My main result concerns who pays cash, and deed party names supply a screen for one source of 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 they carry reduced prices, they 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]. Deleting the same number of unflagged pairs, matched on borough, financing transition, party-record coverage and non-person parties, does not reproduce the decline. Across 999 such draws the median change is +0.44 log points, and no draw is as negative as the screen. Ignoring the 40 most common surnames to limit coincidental matches gives 6.71 [4.69, 8.72]. Because the flag misses relatives with different surnames and can match unrelated people, I read the change it produces as evidence that composition matters, and I would not treat it as the share of the gap caused by family transfers. State transfer reports filed with a random sample of deeds suggest the flag picks up the right transactions. In those reports, 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.

The recorded identity of the cash buyer also matters. When the buyer at the unfinanced sale is a company, the contrast is 1.53 log points [-1.79, 5.44], against 13.77 [10.96, 15.98] when it is an individual or another party. That comparison holds the buyer's type fixed at both sales, so it measures what paying cash adds to the discount companies obtain in any case. Without that control the company-cash contrast is 7.3 log points (Section 5). Applied

to both property types under a common design, the surname screen moves the house contrast from 11.42 to 7.50 log points and the condominium contrast from 0.83 to -0.37 [-1.05, 0.27], so the divergence between property types survives.

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 8 percent for condominiums. Since failure probabilities and continuation values are unobserved, these figures describe what execution risk would have to explain and say nothing about its actual share. An execution-risk mechanism would also predict that the contrast moves with mortgage rates and inventory, and the associations I find along those lines do not survive financing-specific time trends.

To measure that composition, I apply one repeat-sales design to houses and condominiums in the same city and evaluate 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 I cannot see in public records. 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.

I do not observe contracts, contingencies or failed closings, and my financing labels come from record linkage rather than verified loan purpose (Appendix A collects these limits). Nor do I estimate the incidence of the mortgage recording tax.

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). Agreement with that paper's resolved AI-assisted readings does not establish the classifier's accuracy in this house sample. A second companion paper (Loschi 2026b) studies how recorded prices responded to New York's 2019 transfer-tax notches and reports no repeat-sales estimates.

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, and 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, and 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.

The distinction is statutory, and because sellers' prices, financing terms and ownership choices can all respond to the tax, statutory borrower liability is unlikely to equal the ultimate economic burden. Even a cut in a buyer-paid transaction tax is partly passed to sellers: Ke et al. (2025) estimate that buyers kept about 80 percent of the surplus from the United Kingdom's 2020–21 stamp duty holiday. The same boundary creates a measurement problem, since a search of real-property mortgages alone cannot identify co-operative share financing, and ACRIS personal-property filings have to be examined separately.

The proposal considered in May 2026 would have taxed purchases without a mortgage. These records cannot establish its full tax base: an unmatched purchase may involve unobserved borrowing, and the condominium linkage leaves some sales unclassified. The comparison here concerns recorded financing and transaction composition, rather than revenue from that proposal.

[1] Primary sources: New York Tax Law §§253, 253-a, and the NYC Department of Finance mortgage recording tax schedule.

3 Data and record linkage

3.1 Samples and the boundary of replication

The analysis input contains 294,733 sale rows across five boroughs. Excluding Staten Island, whose deeds are recorded by the Richmond County Clerk outside ACRIS, leaves 166,907 house and 86,381 condominium sales. The condominium sample requires an original parcel identifier and a recorded price strictly between \$300,000 and \$6,000,000.

The panel-construction code combines NYC Department of Finance annualized sales with ACRIS real-property master and legal records. Because 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 original input files are preserved byte-identically to repository commit `f404e48320a81e3bfe20127af2ed7c2d9fb5268e` (Loschi 2026c). The house estimates use that panel; the condominium estimates use the rebuilt panel described below.

New public extracts add DOB permit dates, PLUTO geography, ACRIS deed parties and consideration, co-operative transactions and personal-property filings (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). The replication README documents source identifiers, preserved extracts, software versions and the exact computational boundary. The house analysis starts from its archived panel. The condominium panel has been rebuilt from the frozen sale and document extracts, using unit lots 1001–6999. The correction package preserves the earlier outputs and supplies the updated panel, endpoint flags and estimates.

Linkage quality has several dimensions: Bailey et al. (2020) distinguish match rates and sample representativeness from false-link and missed-link errors, but their error rates come from historical census data and are unlikely to carry over to ACRIS. In this paper, agreement with archived flags checks the computation, and the document readings of Appendix H check the financing labels on samples.

3.2 Financing and repeat-sale construction

For houses and condominiums, financing is the recorded 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. I keep the narrow, base and wide labels, and for brevity I call the absence of a matched mortgage "cash", although it can include unobserved borrowing,

missed linkage or instruments outside the rule. Only new mortgages count, so assignments and satisfactions are excluded.

Within each eligible house or condominium parcel, I sort sales by date and keep adjacent observations at least 1,095 days apart, which yields 8,208 house pairs and 6,836 condominium pairs. All sales fall between 2016 and 2025, and 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 eligible, already restricted panel, so a transfer excluded from that panel can still lie between two adjacent sales. 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, and leaves 6,316 pairs. Permit issuance does not establish that work was completed or how much the property's condition changed, so I treat the restriction as a sensitivity check for changes in condition and make no claim that it holds quality constant.

3.3 Parties, consideration and geography

I assign a deed to each house endpoint only when the closest candidate document date within 45 days is unique, and leave ties 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 that the records agree with each other, which is weaker than showing that consideration equals a negotiated retail value.

Grantor and grantee names supply entity, trust/estate and unknown-party indicators, with explicit company forms taking precedence over trust tokens. The narrower institutional rule identifies named financial institutions, and a broader alternative also uses generic lending and servicing terms. I exclude a pair when either party is flagged at either date. These rules pick out institution-associated transactions, which overlap with, but are broader than, legally adjudicated foreclosure, deed-in-lieu or REO status. A four-case instrument pilot, reported in Appendix E, distinguishes a referee foreclosure transfer from a subsequent lender resale and confirms that ordinary DEED metadata can cover foreclosure instruments. 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. Because the identifiers are current, they control for location but cannot capture past building condition. Of 1,227 ZIP-year endpoint cells, 353 have fewer than five endpoints, so fine geographic

adjustment rests on 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

$$y_i = X_i\gamma + Q_i\beta + \varepsilon_i, \qquad \widehat{\pi}_J = \tfrac{1}{2}(\widehat{\beta}_{cf} - \widehat{\beta}_{fc}). \tag{1}$$

Same-financing pairs help estimate the market controls and provide the reference group for the switching coefficients. All fixed effects enter in first differences, and I attach no interpretation to individual coefficients in what is 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

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

it is $\widehat{\pi}_R = c'M_X y$. The joint coefficients instead satisfy $\widehat{\beta} = (Q'M_X Q)^{-1}Q'M_X y$. Since residualizing y leaves the financing indicators untouched, the two estimands can respond differently when market controls change. Table 1 reports both, but only equation (1) gives the financing coefficient in a jointly controlled regression, and I use it for all comparisons.

Neither estimator is causal without further assumptions, because 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 the observed parts of those differences, which falls well short of making financing exogenous.

4.2 What symmetry diagnoses

Suppose the two residual switching means can be written as

$$\mu_{cf} = \pi + u_{cf}, \qquad \mu_{fc} = -\pi + u_{fc}. \tag{3}$$

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 and leaves their difference free, so opposing condition or composition effects can cancel in A while biasing the half-difference.

This point connects the permit and institution analyses: permits address recorded authorisations to change the property, whereas 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 could reflect either that composition or renovation. For that reason I evaluate each restriction through a refitted model with paired uncertainty and do not let either one serve as validation of the other.

4.3 Cumulative evidence and neighbourhood appreciation

In the specification sequence of Table 1, the institutional exclusion increases the permit-free gap, whereas buyer and seller type controls reduce it. The name-based exclusion therefore leaves the positive house contrast in place. I read the change from party controls as descriptive, since the sequence of steps is one of many and does not amount to a causal decomposition.

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 and do not verify 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.

The geography-linked baseline holds the sample fixed before adding local-year controls. 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 it is not the coefficient of financing in the same jointly controlled regression.

I evaluate each change with its covariance in Table 2, where ZIP-year effects change the joint estimate by 0.39 log points, with interval [-1.02, 1.60]. Local appreciation therefore explains no detectable part of the joint house gap, although current geography and sparse cells limit how small a contribution I can rule out.

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.

For every specification I sum observation influences within the original parcel and stack them before computing the covariance, giving excluded parcels zero influence in that specification, which 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, and it confirms the main joint-coefficient comparison. The high-dimensional residual statistic has a shifted resampling distribution, which is another reason to avoid reading its attenuation as an explained component. All intervals are pointwise and leave out 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. As Appendix Table 12 shows, 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.

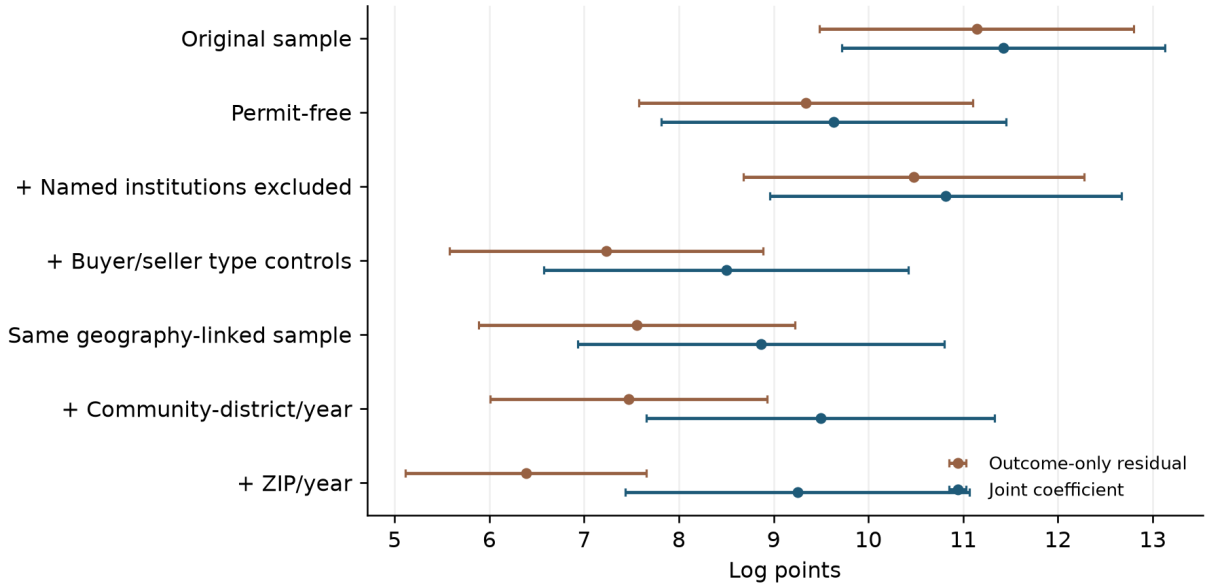


Figure 1: The residual-mean statistic and joint coefficient across controls

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.

Table 3: Outlier-robust estimates of the financing contrast

	Houses, full controls (6,006 pairs)			Borough-quarter only	
Estimator	Contrast	95% CI	Change from LS	Houses	Condos
Least squares (reference)	9.25	[7.17, 11.35]		11.42	0.83
Huber M-estimator	6.73	[5.32, 8.76]	-2.52 [-3.31, -1.14]	8.83	-0.51
Trim growth at 1st/99th pct.	7.66	[5.77, 9.53]	-1.59 [-2.76, -0.45]	9.23	-0.46
Trim growth at 2.5th/97.5th pct.	5.15	[3.55, 6.78]	-4.10 [-5.64, -2.53]	6.91	-0.59

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, so the house contrast remains positive under these estimators even though its magnitude depends on how much weight extreme appreciation receives. Under the common borough-quarter specification, the Huber estimate is 8.83 [7.39, 10.28] for houses and -0.51 [-1.00, 0.02] for condominiums. The condominium contrast is much smaller and moves with the estimator, while the divergence between property types survives.

4.5 Property types and classification sensitivity

When I put houses and condominiums under common borough-quarter controls in Table 4, the condominium estimate is 0.83 log points, with interval [0.06, 1.61]. Even with common controls, the two property types differ in ownership form, in the properties themselves and in matching coverage, so I would not attribute the contrast to standardisation or tax status alone.

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	6,836	1276	1737	0.83	[0.06, 1.61]

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.

The house findings survive narrower and wider financing windows, stable-label restrictions, broader institutional exclusions and consideration-consistent deed matching (Table 5). Disagreement across windows gives a measure of internal consistency: among house sales, strict and base labels differ for 1,020 observations and wide and base labels for 2,193.

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. In a companion paper (Loschi 2026a) I show that such deeds can record an amount set by transfer-tax rules instead

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; Appendix H reports document checks and the remaining uncertainty in financing labels.

of 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, which means that lenders taking title enter the analysis as cash buyers, and 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. This is the direction legal-rule amounts imply, because 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. I treat this exclusion as a sensitivity check. The companion paper's agreement with AI-assisted readings from another sample does not measure error rates here.

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, because a below-market unfinanced sale at the first date raises measured appreciation in a cash-to-financed pair, and one at the second date lowers it in a financed-to-cash pair. Both movements enlarge the half-difference while offsetting 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. I flag a deed when a natural-person grantor and grantee share a surname, using names recorded in the 'surname, given name' form, so company, trust and estate names never match. The flag marks 358 of the 6,006 pairs, and same-surname deeds are 2.5 times as common at unfinanced endpoints (5.7 percent) as at financed ones (2.3 percent). Their prices look like 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, and 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, and ignoring the 40 most frequent surnames in the extract reduces potential 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 probably still contains some related-party transfers.

Removing 358 pairs from 6,006 changes the sample as well as its composition, and the flagged pairs differ from others in location and party records. A matched-deletion exercise asks whether removing any 358 similar pairs would move the contrast as much. Each of 999 draws deletes exactly 358 unflagged pairs, sampled without replacement within cells defined by borough, financing transition, the availability of party records at each endpoint and any entity, trust or unknown party, and refits the full joint model. The design was fixed before fitting and uses no outcome or influence information. The median contrast after deletion is 9.69 log points (a median change of +0.44), the central 95 percent of changes runs from -0.12 to $+0.95$, and no draw falls by as much as the surname screen's -4.38. Adding first-sale and resale eras (2016–2019 against 2020–2025) to the cells gives a median of 9.68 and a central range of changes from -0.09 to $+1.00$. The exercise rules out deletion within these cells as the source of the decline. Whether kinship itself causes it is a separate question, since flagged pairs may also differ in circumstances the records do not show.

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. I drew a random sample of 40 flagged and 20 unflagged cash-endpoint deeds (seed 20260924), read

them from the ACRIS images with AI assistance, cross-checked them against the transfer taxes paid, documented each deed, and then checked every reading against the transfer report myself, agreeing with all 60 (Appendix G). 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, and 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 can omit relationships, and the sample is too small to measure the flag's error rates, so I report these as descriptive rates of disclosed signs.

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

A. Financing contrast by estimator					
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
B. Least-squares contrast by the buyer at the cash sale					
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.

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 pays cash 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 is weak evidence of institutional ownership, since an LLC can represent an individual buyer. The estimated company-cash contrast is smaller, but its interval admits a positive gap as well as zero. The comparison depends on the buyer-type controls in the

nuisance design, which difference out a company indicator at both sales. Companies pay less whether they finance or not, and those controls absorb that general discount. Dropping the buyer-company control, the company-cash contrast is 7.28 log points (parcel bootstrap 95% interval 4.98 to 10.87, 199 draws) against 13.14 for other cash buyers, and without same-surname transfers it is 8.02 (5.63 to 11.97) against 6.35 (4.07 to 8.58). With no party controls at all the figures are 9.44 and 13.60. So company cash buyers do pay less, but by little more than companies pay less when they borrow, and once same-surname transfers are removed the remaining cash contrast is no longer concentrated among individuals. Excluding pairs with same-surname deeds removes almost half of the headline contrast in every specification. Even so, the remaining gap among individuals need not be a premium for certainty, because individual cash purchases also include estate sales, sales to tenants and other circumstances that names do not reveal.

5.3 Condominiums and a shared-address screen

Condominium endpoint deeds are matched on unit BBL, date and amount in the preserved ACRIS extracts. The party snapshot retains the earlier extract and adds records for newly required deeds. A unique deed is linked for 13,639 of 13,672 condominium endpoints, and 33 ambiguous matches are left unlinked. The same surname rule flags 195 of 6,836 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], while the condominium contrast moves from 0.83 (bootstrap interval [0.05, 1.63]) to -0.37 [-1.05, 0.27], a paired change of -1.20 [-1.65, -0.78]. Ignoring the 40 most common condominium surnames gives -0.18 [-0.87, 0.49]. 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] cannot be distinguished 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. I therefore treat this screen as a broad sensitivity check that says little about kinship.

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

A. Common borough-quarter design (Table 4)				
	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	6,836	0.83 [0.05, 1.63]	195 -0.37 [-1.05, 0.27]	-1.20 [-1.65, -0.78]
B. Headline house specification: excluding pairs with				
	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 13,639 of 13,672 endpoints (33 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

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

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

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

This inversion makes the required loss explicit while leaving q as an assumption. I do not substitute national mortgage denial rates for q, because a denied application differs from the failure of an accepted purchase contract.

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,

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

Failure probability q	Houses: $100d^*$	Condos: $100d^*$
0.1	87.2	7.5
0.2	38.8	3.3
0.3	22.6	2.0

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; contract-failure probabilities are unobserved. No common relisting technology is imposed as an empirical fact.

or 63 percent at the Huber estimate, compared with about 8 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. That is still several times the condominium figure, and the screened condominium contrast requires no loss at all. A seller whose financed buyer fails still owns the house and can relist it. Loss aversion makes relisting slower for sellers facing a nominal loss (Genesove and Mayer 2001), but that evidence concerns asking prices and time on the market and involves far smaller losses. A loss of a few percent is in line with ordinary carrying and relisting costs. A loss of more than half the property's value shows how much an execution-risk explanation asks of the house contrast. These calculations state what the model would require, and estimating the mechanism's actual share would need contract failure rates and continuation values that I do not observe.

Two qualifications apply. First, equation (5) evaluates the counterfactual in which the entire observed contrast represents a causal financing premium, which my design does not establish, and so shows what execution risk would have to explain under that assumption. Second, failure probabilities could differ between the two property types, and 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 testing it would require contract-level data.

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, and in this paper 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 net of the general company discount, 1.53 log points, is below both. Part of a contrast estimated from public records may therefore reflect transfers that vendor data would remove before estimation. The samples, periods and estimators differ, and the comparison says nothing about whether the remaining gap equals an execution-risk premium.

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. I therefore use the co-operative comparison as a measurement sensitivity check, with the full assignment and adjacency audit in Appendix D, and it cannot serve as 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. Because quarterly effects do not absorb all within-quarter weekly-rate variation, I also include separate uninteracted rates at both endpoints. The object of interest is the financing-specific association with rates.

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, and the two rate interactions cannot be jointly distinguished from zero in the calendar-dependence sensitivity (Appendix Table 11). Rate movements, time trends and financing composition are hard to disentangle in these data, so the rate evidence does not identify the execution-risk mechanism. The PMMS methodology change on 17 November 2022 adds a further measurement caveat. The augmented model's reference contrast is evaluated at

4 percent and at the trend origin, and it should not be read as an unconditional average financing gap.

Table 9: Credit conditions at both sale dates

Financing × 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 fails to reject zero endpoint interactions (Appendix F reports both specifications). As with rates, this check is unable to isolate a liquidity or execution-risk mechanism.

7 Conclusion

I find that the financing contrast in these repeat sales falls from 9.25 to 4.87 log points when I exclude pairs with a same-surname deed, while deleting comparable unflagged pairs has little effect. Transfer reports point to unusual conditions in some flagged transactions, but shared surnames remain an imperfect relationship screen, so I read the change as a measure of sensitivity to that screen. The share of the original gap caused by transfers between relatives remains unknown.

Once I control for buyer and seller type, the company-cash contrast is 1.53 log points, with an interval of [-1.79, 5.44]. This imprecise conditional association leaves open whether the certainty premium is zero, and a company name says little about the buyer's investment purpose. Differences in sale circumstances and property condition that public records miss may also affect the contrast.

Financing measurement is a separate limit. In a new probability sample, AI-assisted document readings supported purchase borrowing for two endpoints with no matched mortgage, and eighteen such endpoints remained unresolved. Correcting the two documented labels

barely changes the estimate, but the accuracy of the unresolved labels is unknown, and errors concentrated on particular pairs could matter more than random recoding suggests.

For policy, the observed cash discount provides an uncertain basis for setting a tax intended to offset a cash buyer's closing advantage. Financed buyers already face a statutory mortgage recording charge, although its ultimate burden can pass through prices and loan terms. Identifying the value of certainty requires evidence on offers, contingencies and failed closings, or financing variation unrelated to transaction composition.

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). The master script reproduces the house analysis from its preserved inputs. The corrected condominium panel, derived surname flags and offline rerun script are supplied in the accompanying correction package. The dataset release is version 1.2.0, doi:10.5281/zenodo.23557585; the replication README distinguishes the archived house inputs from the corrected condominium inputs. The companion papers' replication packages are deposited separately (Loschi 2026b,a); they address separate questions using overlapping public records, and the deed-classifier exclusion uses the companion classifier as described in Section 4.5.

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.

A Measurement limits

I measure every transaction attribute from records. The document checks cover only a small sample of deeds and financing labels, and I have not checked the rest against the source instruments. A financed label means a purchase-window mortgage was matched and a cash label means none was, and neither confirms how the purchase was paid for. A shared surname or mailing address marks a possible relationship, missing relatives with different names and sometimes matching unrelated people. Transfer reports in a small random sample of deeds show non-market signs in 40 percent of flagged deeds and 10 percent of other cash deeds, but the declarations are self-reported. Permit issuance shows authorised work that may not have been completed, and institutional-name rules mix foreclosure transfers with lender resales. Financing arrangements can also make a recorded price exceed the net price: in Cook County deed records, Ben-David (2011) finds that some highly leveraged buyers inflated recorded prices to draw larger mortgages, which would raise financed prices relative to cash ones. The public records used here do not show seller credits or concessions, so I cannot measure that channel. 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, and I recompute everything downstream of it.

I checked the financing labels on two samples. A review of 53 pairs confirms 54 of 56 financed labels and finds a purchase mortgage behind 4 of 50 cash labels, and dropping the 4 affected pairs leaves the contrast at 9.40 log points (Appendix H). Under random errors in either direction, at the rates used in the sensitivity exercise, the average contrast stays between 5.82 and 8.96 log points and the surname screen still removes about half of it. Errors related to price growth matter more, since the contrast reaches zero if the recoded labels sit 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.

B 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

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

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

$$\mathrm{IF}_{\beta,i} = (R'R)^{-1} R_i' e_i. \tag{7}$$

I apply the half-difference contrast, sum within parcel and stack across specifications. The covariance is the cross-product of these stacked cluster influences (CR0), and paired differences use the corresponding covariance terms.

A numerical check perturbs five parcel weights in the full ZIP-year model, refits the nuisance parameters and matches both analytic derivatives to finite differences within 10^{-7}. The common bootstrap draws multinomial weights over the 8,127 original house parcels (seed 20260909) and refits the original, permit-free, linked-party-control and ZIP-year specifications in each of 999 draws. Redundant nuisance columns are projected out. Intervals are pointwise, with no multiple-testing correction, and they condition on the observed linkage and classifications.

C Replication map and classification definitions

The master script `code/reproduce.py` parses the preserved rate extracts, rebuilds the enriched pairs, estimates the house models and bootstrap, audits co-operative assignments, estimates the outlier-robust models, runs the numerical checks and regenerates the exhibits. The replication README gives setup instructions, package versions and a file-level exhibit map, and keeps archived inputs separate from newly retrieved extracts and derived outputs.

Table 1, Table 2 and Figure 1 come from the repeat-sale results and joint bootstrap outputs. Table 4 uses the archived house panel and the corrected condominium panel described in Section 3.1. 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, and Table 8 is computed from the displayed formula. The matched-deletion exercise of Section 5 is produced by `run_matched_deletion.py`, which records its frozen plan, seeds, deletion masks and source hashes.

The house and condominium loan windows are document-dated purchase windows, whereas the co-operative windows use personal-property recording dates, so the two differ in coverage and timing. Party classifications are reproducible regular-expression categories, with unknown party records kept in their own category instead of being counted as individuals. The release keeps the audit fields and review queues, so later adjudication can be versioned without changing the underlying rules.

D Co-operative assignment and adjacency audit

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

A building-level UCC filing need not be a particular buyer's purchase loan. Several sales and filings can fall close together, some sales lack a recovered unit identifier, and filings can concern other credit. When I count possible competitors for a filing, the co-operative extract therefore includes every observed price and incomplete-unit record as well as the estimation sample.

The broadest rule, the expanded count rule, counts eligible sales within $d \pm 45$ days of a sale at date d 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, and the rest are ambiguous. The expansion matters. Of 60,496 purchases labelled financed, 4,385 have no filing in the base $[-15, +90]$-day window, and counting all competing sales also changes assignments.

The alternative rules require exactly one candidate filing and exactly one observed sale eligible to claim it, with narrow, base and wide windows of $[-15, +30]$, $[-15, +90]$ and $[-15, +180]$ days. Documents linked to several parcels are excluded. These conditions stop one filing from being used twice but say nothing about its debtor, collateral unit or purpose, and zero candidates means only that no record was matched.

Adjacency uses all observed known-unit sales before endpoints with ambiguous financing or out-of-band prices are dropped. Forming pairs within the estimation sample would bridge an intervening unit sale in 271 cases, so every row of Table 10 forms adjacency on all known-unit sales. Base timing uniqueness leaves 1,765 pairs in 1,082 buildings, and its interval includes zero, as do the narrow and wide alternatives. That is consistent with a small co-operative contrast in selected, differently measured samples, and it is silent on whether the mortgage recording tax causes the house gap.

Two timing-unique candidates were inspected. One filing names the recorded sale unit, and the other pledges a different unit in the same building. The second sale forms no repeat pair, so rejecting its link leaves Table 10 unchanged, and the first still lacks confirmation of purchase purpose. The pair shows how timing-unique assignment can fail, although two cases cannot give an error rate.

E Targeted instrument pilot

On 10 September 2026 four queued cases were read through the public ACRIS image viewer, with AI assistance: one broad-only house flag, one specific-institution house flag and two timing-unique co-operative assignments. Page images, document identifiers, notes and source

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 and does not prove an unlevered purchase. Estimates use the September 2026 co-operative extract.

hashes are in the replication package. The earlier 90-house and 90-co-operative queues were not completed and are preserved as they stand; their sampling weights do not apply to these four cases. The focused review of Appendix H replaced the planned two-human-reader review.

House deed 2023051600851001 recites the lender's acquisition under referee deed 2019020500740001. The earlier instrument describes a mortgage foreclosure and conveyance to the lender, and the later sale reads as a lender resale, which the broad name rule flags and the narrow rule misses. House deed 2019093000695001 is a referee's deed in foreclosure and is caught by the narrow institutional-buyer rule. Both cover records carry DEED metadata, and the two cases separate mechanisms within the institutional flags.

For co-operative filing 2025020500312001, the sale unit and the collateral are both 7AB, but the loan's purpose is unverified. Filing 2020020400882001 pledges unit 4E, while its timing-assigned sale is for 1CD. Since that unit has one observation in the sale extract, the false link affects no repeat pair. I reject it as evidence on that sale's financing and leave the sale's label unchanged. A timing-only rule cannot catch this kind of mismatch even when no competing sale exists.

F Historical market conditions and shared calendar shocks

The inventory extension uses a single indicator, the log base two of StreetEasy's prior-month borough sales inventory in its Single Family category, centred on each borough's mean log inventory over December 2015–November 2025, so a unit is a doubling. The September

2026 vintage covers both endpoints of all 6,006 pairs. Listing coverage need not match the administrative one-to-three-family sample, and inventory reflects demand, supply and platform participation as well as anything related to closing risk. Because of revisions and unknown publication dates, a prior-month value may not have been public 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 last keep fixed geographic differences in financing out of the inventory terms). I then add financing-specific linear calendar trends as a sensitivity check. I wrote this plan after checking coverage and before fitting, without registering it externally, and report both specifications.

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 without exact small-cluster inference.

Parcel-only uncertainty is in the replication output, and the table also allows dependence whenever two pairs share a parcel or either endpoint calendar period, whichever sale the period belongs to, counting overlapping memberships once. The calendar-quarter calculation has 40 nodes and the calendar-year version ten. Negative eigenvalues in the covariance are set to zero and their magnitudes reported. The t and F references are approximate, intervals are pointwise, and dependence across non-overlapping periods is left out.

With trends, the inventory interactions are jointly indistinguishable from zero under either calendar aggregation, and the credit exercise gives the same answer once common weekly rates are included. The associations are sensitive to the time specification, which is weaker than showing that market conditions are irrelevant. The headline contrast comes from the main model and should not be confused with the reference contrast of either interaction model.

G Deed-image review of the same-surname flag

A seeded script drew the sample before any deed was read, taking 40 endpoint deeds at random among flagged pairs and 20 among unflagged cash endpoints of the headline sample (seed 20260924). The 60 recorded documents run to 684 pages, retrieved from the public ACRIS image service, and each has a cover page, the deed and a New York State real property transfer report (RP-5217NYC). Item 12 of the report gives the full sale price and item 14 lists conditions of the transfer, among them (A) a sale between relatives or former relatives, (C) a buyer who is also a seller and (E) a deed other than a warranty or bargain-and-sale deed. Prices and conditions come from the report, while parties' addresses and the city and state transfer taxes, whose rates imply the taxed consideration, come from the cover pages. Where the report price and the taxed consideration disagree, I record both and class the case as unclear. The readings were made with AI assistance and documented deed by deed. I then checked every one of the 60 against crops of the transfer report, comparing the price, the conditions declared and the buyer and seller names, and agreed with all 60. The tax amounts on the cover pages were not rechecked. Readings, my answers, image crops and scripts are in `audit/deed_image_review`.

A deed shows a non-market sign when it declares condition A, C or E, or reports a zero price with no transfer tax. The rule is conservative because conditions are self-declared. A transfer from a father to a son of the same name, for example, may declare none. Among flagged deeds, 9 more list a buyer among the sellers or give one address for both sides 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 supports the flag as a marker of non-market transfers, but it is too small to measure the flag's sensitivity or specificity.

H Focused validation of the house contrast

A first pass checks the archived rows behind the 6,006 pairs, which cover 5,957 parcels and 11,963 distinct sale endpoints. Every pair's price change and financing flags reconstruct from those rows, the three financing windows are nested, and the linear coefficient weights reproduce the joint contrast. This confirms the computation, though it leaves the archived linkage untested.

Ten pairs have an endpoint with more than one house record on the same parcel and date, and three financed endpoints have no positive mortgage amount. Neither proves a wrong label, and dropping these cases and applying the existing safeguards together leaves

5,579 pairs and a contrast of 9.02 log points. The archived panel does not keep the mortgage document identifiers needed to verify loan purpose directly.

Table 12 reports the full refitted sequence.

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. The influence exercises remove whole parcels. Absolute- and positive-influence exclusions are selected using the fitted outcome; they are diagnostics and should not be used as 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 restrictions check internal consistency and do not validate the classifications. The complete model outputs and nominal parcel intervals are retained in the replication files.

Influence comes from the fitted joint model, aggregated by parcel, and every deletion refits all nuisance controls. Removing the single most influential parcel leaves 9.07 log points, and removing the 60 with the largest absolute influence leaves 8.51, so no single record drives the estimate. Removing the 60 largest positive influences lowers it to 4.68. Because that deletion selects on the outcome, Section 4.4 handles the same concern with robust estimators at conventional tuning constants. Omitting each borough and each of the three largest ZIP

groups is reported in full.

The replication output also reports adversarial perturbations to pair growth with labels and design held fixed, moving the outcomes of the parcels with the largest coefficient weights against the sign of those weights. These are exact linear sensitivity calculations for contaminated outcomes and bound neither misclassification nor unobserved renovation.

Document review. A price-blinded queue drew 20 pairs at random from each financing-switch arm and added 13 targeted cases (the pairs on the ten most influential parcels, five per arm, and the three pairs with a financed endpoint lacking a positive mortgage amount), giving 53 pairs and 106 sale endpoints. The protocol called for two human readers, but the readings were made with AI instead. ChatGPT read the deed and every candidate mortgage at each endpoint, recording borrower, collateral, purpose and pages, and Claude checked each reading against ACRIS open data (parties, amounts and every mortgage on the parcel from 15 days before to 90 days after the sale) and sent inconsistent readings back. Readings were locked before the analyst key was opened. Claude did not read the images, and the open-data check uses the same public records the labels came from, so the rates below measure agreement with these readings.

The documents confirm a purchase mortgage at 54 of the 56 financed endpoints. Of the other two, one shows a new loan on a transfer in which a buyer is also a seller and the other shows land loans to the buying company. Both are conflicts, and neither indicates cash. Among the 50 cash endpoints, 4 have a same-day purchase mortgage in the documents (8 percent). The pooled count mixes random and targeted cases, so it has no population interval. Two of the errors come from the city's index dates. In one pair the mortgage, and in the other the deed, is indexed a year from its true date, so the window missed the loan. The other two are loans recorded against two adjoining lots. Since a missing loan cannot confirm a cash purchase, pair-level error bounds stay wide, and in the two random arms between 6.2 and 100 percent of pairs carry a label error, and all 3 observed errors are in the financed-to-cash arm.

Dropping the 4 pairs with a documented error 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 a loan recorded 25 days after the transfer and a two-lot development loan as unresolved leaves 3 documented errors and 9.39 log points. Hidden purchase mortgages can push the contrast either way, depending on which pairs carry them. Correcting the documented errors raises it slightly, and the random recoding below lowers it slightly. Readings, the open-data check and the scoring code are in `audit/cross_ai_review/` and `code/cross_ai_review.py`.

Errors outside the review. Table 13 asks how far the contrast would move if similar errors existed among unreviewed pairs. It starts from the corrected sample (9.40 log points, or 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, among the 712 unreviewed financed-to-cash pairs and the 1,019 unreviewed cash-to-financed pairs. The shares are scenarios set by what the review found, namely 2 or 3 of 20 sampled financed-to-cash pairs, depending on whether the two-lot loan counts, and none of 20 cash-to-financed pairs, up to Wilson upper limits of 36 and 16 percent. I fit each draw 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 are distributions of point estimates under assumed recoding and should not be read as confidence intervals. Table 14 relaxes the three restrictions of this exercise.

Both directions and the whole sample. Panel A recodes unreviewed labels anywhere and in either direction, moving a share of the 2,613 cash endpoints to financed, at 8 or 18.8 percent (the pooled reviewed rate and its nominal Wilson upper limit), and a share of the 9,293 financed endpoints to cash, at 3.6 percent (the two unresolved financed readings counted as errors) or 6.4 percent (the Wilson upper limit for no error in 56). Because the pooled counts include targeted cases, these are scenario settings. Cash–cash and financed–financed pairs can now become switching pairs, and the arms are rebuilt. Random recoding mixes pairs of ordinary growth into both arms and pulls the contrast toward zero, to between 5.82 and 8.96 log points on average, without changing its sign, and the surname screen still removes 47–51 percent of it.

Documented failure modes. Panel B follows the two mechanisms the review found. An index-date error or a late recording appears as a mortgage in the wide window but not in the base window, and recoding every such unreviewed cash endpoint (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. A recoded label raises or lowers the contrast depending on the pair's growth net of time and location effects. A hidden mortgage in a weak-growth financed-to-cash pair, for example, removes it 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 summarises that tilt as the mean standardized effect of the chosen labels, zero when errors are unrelated to growth. The contrast falls steadily and reaches zero at 0.70 standard deviations, or 0.45 without same-surname pairs. Choosing the labels with the largest effect

in either direction bounds the estimate between −15.70 and 25.49 log points. The bounds are wide because a few hundred labels can be placed on the pairs that matter most, and they mark the limit of what the data can exclude. The failure modes in Panel B show no such tilt, but the review is too small to rule one out.

Table 15 applies the documented corrections and the screens together on the 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. The contrast and the share removed by the surname screen hold under random errors in both directions and under the error patterns the review found. They would fail if errors sat on the pairs whose growth works against the contrast, which the review can neither detect nor exclude.

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 H); 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 of point estimates. 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), which are scenario settings and should not be read as measured error rates. `code/financing_stress_test.py`.

A targeted document follow-up. A later AI-assisted pass, unblinded and only partly checked by me, read the deeds, transfer reports and candidate mortgages of seven influential endpoints picked out by a tabular trace of all 11,963 sale endpoints. It found two problems.

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%
A. Random, both directions, all pairs					
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%
B. Documented failure modes					
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%
C. Related to growth, hidden 8.0%, false financed 3.6%					
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. Ranges are scenario ranges of point estimates. `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 4.5, Appendix H). `code/misclassification_checks.py`.

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 the transfer report marks as commercial. With no allocation to the house in the pages read, these pairs do not compare prices of one house. Four cash endpoints have a mortgage indexed one or two years before the dates written in the mortgage and note, which match the buyer's deed, so the window missed them, as it missed the two index-date errors above. Purchase financing is inferred from the same-day deed, mortgage and collateral, without a use-of-proceeds clause. For case P02, reopening the referenced instrument showed that document 2022042100383001 is a power of attorney. The endpoint deed is 2022042100383002, and the operative mortgage and note are dated 7 April 2022, while the mortgage cover indexes 7 April 2020 and cross-references the power of attorney. The date-rule recodings remain sensitivity scenarios. Findings are in `audit/targeted_followup`.

Because the seven cases were chosen for influence, correcting only them could favour one direction, so Table 16 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 from correcting the seven cases alone, and slightly above the headline. 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, because it requires the loan to be recorded near the sale, to name the buyer and to cover only the purchased parcel. `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]
A. Seven targeted cases			
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]
B. The same rules applied to every pair			
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.

H.1 A fresh probability sample of financing labels

On 9 October 2026 a seeded draw sampled endpoints not covered by earlier reviews. Excluding 111 reviewed endpoints left 11,852, of which 2,598 had no matched mortgage and 9,254 had one. The draw took 20 from each stratum without replacement and hid the recorded labels in an analyst key. The queue, key and input hashes were frozen, and readings were locked before the key was opened.

A frozen parcel-and-date search found candidate instruments for 22 endpoints, and AI readings examined the operative deeds and mortgages for all 22. Ten received a second AI reading, including the three cases first supported only by recording covers and those raising seller financing, date, parcel or transfer-type questions. No human reader repeated them. One instrument states purchase money outright. The other purchase connections are inferred from the acquisition, the new loan, the buyer–borrower match, the collateral and the timing. The remaining 18 endpoints had no candidate in the initial search and stay unresolved, since a failed search is not evidence of a cash purchase.

After unlocking, all 20 financed endpoints had supporting mortgage evidence, and so did 2 of the 20 unfinanced ones. One mortgage carries a January 2022 index date but a January 2023 operative date matching the acquisition. The other secures the house together with a separate parking parcel. Both recodings had a second reading before unlocking, and the date conflict and the bundled price are kept as separate issues.

Letting the unresolved cases agree or disagree gives a design-weighted bound of 2.19–21.92 percent for label discrepancy among previously unreviewed endpoints. Exact finite-population hypergeometric inversion within each stratum, with a Bonferroni split across four one-sided tails, gives a conservative 95 percent envelope of 0.19–37.26 percent. Both condition on the AI readings being right and cannot account for correlated reading errors. The envelope is too wide to pin down a population error rate.

Table 17 corrects only the two sampled labels, keeping the original sample and model. The contrast moves from 9.25 to 9.23 log points. Excluding the house-and-parking bundle gives 9.26, and also excluding the sampled lender resale and referee auction gives 9.28. Intervals are conditional analytic parcel-clustered intervals, and no unsampled label is extrapolated. Readings, design, comparison and code are in `audit/probability_sample_20261009/`, and `second_read/analyze_probability_sample.py` reproduces the calculation.

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

A. Twenty endpoints from each recorded-financing stratum			
Recorded label	Population	Supported mortgage	Unresolved
Without matched mortgage	2,598	2	18
With matched mortgage	9,254	20	0
B. Sensitivity in the original headline model			
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 and are 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 and does not imply that their prices are unusable. Archived data and headline results remain unchanged.

Source checks of the unresolved endpoints. A subsequent, unblinded AI-assisted review added page-specific observations for all 18 unresolved endpoints. It read 17 transfer-report first pages and four complete operative deeds; the other deeds remain partial. Four loan instruments were inspected across three endpoints. At N002 a home-equity credit line is dated 302 days after the deed; at N010 the earlier mortgage names the sellers as borrowers, and the complete deed contains no express assumption of it; at N011 the buyer's building and land loans are dated one year after the transfer. None establishes financing of the sampled purchase. All 18 endpoints therefore remain unresolved, and the discrepancy bounds above are unchanged.

The deeds also identify two foreclosure conveyances (N005 and N016), and two transfer reports declare a sale between relatives or former relatives (N013 and N027). These observations describe the recorded transactions; they do not establish an absence of borrowing. Case notes keep the deed recital, the transfer report's full sale price and the ACRIS index amount separate. No new classification-error estimate or regression recoding is based on this follow-up. Page citations, retained images and remaining evidence gaps are in `audit/unresolved_followup_20261023/`.

Declaration of generative AI and AI-assisted technologies in the writing process

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 H 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 H.1 was read with Codex, with a second AI instrument reading for ten cases; it has no independent human validation. The 60 transfer-report readings of Appendix G were made with AI assistance and then checked one by one by the author. The research uses administrative public records; no human subjects were involved.

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.

Itzhak Ben-David. Financial constraints and inflated home prices during the real estate boom. *American Economic Journal: Applied Economics*, 3(3):55–87, 2011. doi: 10.1257/app.3.3.55.

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.

David Genesove and Christopher Mayer. Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics*, 116(4):1233–1260, 2001.

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.

Qiulin Ke, Bing Zhu, Michael White, and Bin Chi. Transactions tax change during the pandemic: A study of the UK housing market. *Journal of Housing Economics*, 69:102088, 2025. doi: 10.1016/j.jhe.2025.102088.

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: Round prices and bunching at New York City's 2019 transfer-tax notches. 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.

Office of the New York City Comptroller. Comments on New York City's fiscal year 2027 adopted budget. 12 August 2026, 2026. URL https://comptroller.nyc.gov/reports/comments-on-new-york-citys-fiscal-year-2027-adopted-budget/.

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.

Geoffrey K. Turnbull and Arno J. van der Vlist. The price of ignorance: Foreclosures, uninformed buyers and house prices. *Journal of Housing Economics*, 57:101844, 2022. doi: 10.1016/j.jhe.2022.101844.