Drawing the Line: The Spatial Non-Neutrality of Land Value Taxation*

Pedro Aldighieri†

July 2026

Abstract

Land value taxes (LVTs) tax land values, exempting buildings. Arnott and Stiglitz (1979) show that a 100% LVT can finance local public goods without distortion. This result fails once assessed land values capitalize neighboring spillovers, which enter the base only when they cross property lines. An LVT then favors merging parcels under positive spillovers and splitting under negative ones, and subsidizes externalities that depress neighbors' assessed land, functioning as an inverse Pigouvian subsidy. Difference-in-differences estimates from Cook County splits and mergers (2001–2023) confirm the required mechanism: assessed land values move in opposite directions, reversing between nuisance and positive-spillover exposures.

Keywords: land value tax, property tax, externalities, land assembly, urban economics, spatial economics, Georgism

JEL Codes: H21, H23, H71, R32.

*I thank Netanel Ben-Porath, Julien Manilli, Therese McGuire, and Joan Youngman for the helpful comments and suggestions.

†Department of Economics, Northwestern University. E-mail: pedro.aldighieri@u.northwestern.edu

1 Introduction

George (1879) famously argued that taxes on land rents need not create deadweight loss because land is fixed, motivating his proposal for a "single tax." A land value tax (LVT) implements this idea by taxing a parcel's value as if vacant, excluding the value of structures and improvements on the parcel. Stiglitz (1977) and Arnott and Stiglitz (1979) formalize this intuition: aggregate differential land rents at the efficient city size equal expenditures on local public goods, so a 100% land tax on those rents can finance the public goods without distorting private choices. This Henry George Theorem, often called the "golden rule" of local public finance, led economists to put forth LVTs as an ideal non-distortionary tax (Vickrey, 2001; Mirrlees et al., 2011; Oates and Schwab, 2009).

This paper shows that this benchmark can fail because land values reflect spillovers from neighboring properties. Even though land itself is inelastic, the LVT base depends on how parcels are partitioned, and owners can manipulate it through partition choice.[1]

In practice, a land value tax does not tax some pure location rent. It taxes a bundle of bare site rent and spillovers at every scale – a day-care across the street, a nearby park, a citywide transit line. In cities, local spillovers between land parcels are substantial. Retail amenities raise nearby property values by 2–3% (Pope and Pope, 2015), while crime reduces them by 4% (Linden and Rockoff, 2008); environmental nuisances can be even larger – homes within 1–1.5 km of shale gas development lose 10–17% in value (Muehlenbachs et al., 2015), and brownfield cleanups raise nearby property values by about 5–12% (Haninger et al., 2017). The value of "land" therefore depends on what surrounds the property, and crucially on where boundaries are drawn.

In this paper, I build a model showing that when spillovers between neighboring parcels are capitalized into land values, an LVT can subsidize developments that harm neighboring parcels, inverting the logic of Pigouvian taxation. If two improvements lie within a single assessment unit, the assessment counterfactual is the absence of *both*, so any spillover between them disappears from the tax base. If each lies in a separate unit, the counterfactual for each land assessment removes only its own improvement *while keeping the other standing*. The spillover from one improvement remains in the other parcel's

[1]Despite this theoretical appeal and recent policy interest, pure LVTs remain largely unimplemented. U.S. experience is mostly with split-rate property taxes that tax land at a higher rate than improvements. The main U.S. examples include Allentown and Harrisburg, Pennsylvania, which tax land values at rates between 2.45% and 3.1%. For a survey of international experiences with LVTs, see Hughes et al. (2018), chapter 6. Estonia is a close national implementation because it levies an annual tax on land value rather than buildings, but the rates are between 0.1–2.0% and homeowner relief is substantial. OECD Revenue Statistics 2025 reports Estonian taxes on property at 0.5% of total tax revenues in 2023.

assessed land value, and is taxed. The LVT base therefore depends on where boundaries are drawn. And boundaries are themselves endogenous.

Two distortions follow as a result. The first concerns how parcels are grouped into ownership units. Since only cross-boundary spillovers enter the tax base, owners can shrink their bill by merging parcels when the spillover between them is positive and splitting when it is negative. Holding fixed what is built, this regrouping only relabels which property each spillover is booked to. The same structures stand and only revenue moves, as the base erodes toward bare land value and the differential rents the Henry George Theorem relies on disappear. This becomes a real inefficiency when the grouping itself matters, either because boundaries constrain what can be built or because ownership structure governs how it is financed and operated. Tax minimization then pulls the configuration away from its value-maximizing form.

The second distortion concerns what gets built, and it matters even when boundaries do not constrain developments. An LVT rewards any development whose negative spillover crosses a property line, because exposing the spillover shrinks the assessed land base and refunds a share of the damage. The owner builds the nuisance whenever that refund exceeds the private value she loses to it, including the capitalization into parcels she controls. This is an inverse Pigouvian subsidy that tips developments from rejected to built. Both distortions require that a single decision-maker bear the tax liability across the parcels involved, whether through common ownership, Coasean bargaining with transfers, or, crucially, a split that creates the cross-boundary exposure in the first place.

This problem is specific to LVTs. The standard property tax applies a single rate to the combined value of land and buildings, so whether parcels are merged or split, the total assessed value is the same. Only its decomposition into "land" and "buildings" shifts.[2]

Most parcels have buildings on them, so bare land value is rarely observed directly. Any LVT requires assessors to estimate what the land alone would be worth. The empirical section tests whether, in practice, those estimates behave as the model predicts. The model generates a simple pattern dictated by the sign of the spillover: a positive spillover should raise the assessed land value at a split and lower it at a merger; a negative spillover does the opposite.

Whether assessments actually behave this way is not obvious from assessment manuals alone. Standard practice instructs assessors to value land at market value (International

[2]This distinction becomes especially important because LVT proposals often contemplate rates of 5–20%, far exceeding typical property tax rates of 1–2%, which would magnify any distortions from boundary effects.

Association of Assessing Officers, 2017), but assessors might price spillovers into the building component rather than the land component, work at spatial units too coarse to register parcel-specific shocks, or update values slowly relative to boundary changes.

To test these patterns, I use 7,330 simple splits (one parcel becoming two) and mergers (two parcels becoming one) in Cook County, Illinois between 2001 and 2023. Cook County is not an LVT jurisdiction, so the outcome is the assessor's annual decomposition of assessed value into land and improvements – the bare-land component any operational LVT would have to impute. Using a difference-in-differences design, I find that the estimates match the model's predictions: nuisance and positive-spillover events move in opposite directions, and raw split and merger estimates reverse signs within each spillover class. The estimates imply annual tax bill changes on the order of 14–28% of the affected land base times the tax rate. On valuable urban footprints, this recurring cost can exceed the one-time administrative cost of a boundary reconfiguration.[3]

Common ownership is also empirically relevant in the sample. Over half of mergers combine parcels already held by the same owner, more than a third of splits leave both child parcels with the same owner one year later, and roughly a quarter of Cook County parcels belong to owner groups holding multiple parcels on the same tax block.

The mechanism proposed not only has empirical relevance, but also reflects a broader principle. When tax assessment boundaries do not match externality boundaries, taxes on the resulting "cut" create distortions. Just as value-added taxes restore efficiency by eliminating the tax cascade at firm boundaries (Xing et al., 2024; Hoseini and Briand, 2020), land value taxes would require boundary alignment or spillover credits to maintain efficiency. The boundaries in our setting are spatial, but the same logic applies wherever taxable value is composed across notional boundaries that taxpayers can redraw, like firms, entities, or jurisdictions.

Related literature. Several papers document the capitalization of property taxes, public goods, amenities, and nuisances into property values, beginning with Oates (1969). Causal estimates consistently find capitalization of school and property taxes (Giertz et al., 2021; Lyu, 2024), public investments (Cellini et al., 2010), environmental nuisances and health risks (Currie et al., 2015; Muehlenbachs et al., 2015), and local amenities (Pope

[3]To illustrate, Cook County charges $100 for a division or consolidation petition covering up to three tracts (Cook County Assessor's Office, 2026b); Chicago lists a $2,000 subdivision application fee, and one practitioner estimate places total Chicago subdivision costs at roughly $10,000–$50,000 (City of Chicago Department of Transportation, 2024; Birchwood Law, 2025). A reconfiguration costing $10,000 pays for itself within a year on land bases above roughly $180,000–$360,000 under a 20% LVT.

and Pope, 2015; Linden and Rockoff, 2008; Haninger et al., 2017). This paper shows theoretically that, under an LVT, capitalization generates a new distortion channel through assessment-boundary endogeneity, and shows that empirically assessed land values respond to boundary changes in the direction the model predicts.

This paper also adds to the literature on the efficiency of land value taxation and the robustness of the Henry George Theorem. I am not the first to note that the Arnott–Stiglitz benchmark can fail. A classic exchange established that land value taxes are non-neutral when assessed values capitalize a parcel's own future development (Bentick, 1979; Mills, 1981), and that neutrality returns under pure site-value assessment (Tideman, 1982; Wildasin, 1982). That channel runs through time. The channel here runs through space, and site-value assessment does not close it, because valuing a site as if vacant still requires deciding which neighboring improvements remain standing, a decision that moves with the boundaries themselves. Existing U.S. evidence studies the real effects of land taxation on development (e.g., Banzhaf and Lavery, 2010), but this paper instead shows that the operational LVT base can itself become a choice variable, a failure mode internal to the canonical environment rather than an external distortion introduced into it.

2 The Land Taxation Model

Parcels, partitions, and developments. Let $N = \{1, \ldots, n\}$ denote primitive land parcels. A partition P of N groups parcels into *blocks* $b \in P$: taxable assessment units, which in practice coincide with separately deeded and ownable tax parcels. A *development* z specifies what is physically built across the footprint. Both feasibility and costs may depend on the partition – a mall is hard to operate across separately owned titles, and a row of separately financed homes is hard to hold as one. Write $\mathcal{Z}(P)$ for the developments feasible under P and $c(z, P)$ for development and ownership costs.[4] Boundaries are *technologically neutral* when $\mathcal{Z}(P) = \mathcal{Z}$ and $c(z, P) = c(z)$ for every P; redrawing them is then pure paperwork.

Values. All values are present-value stocks. Parcel i has baseline unimproved land value ℓ_i, with $\bar{\ell} \equiv \sum_{i \in N} \ell_i$. A development z creates aggregate nonland value $s(z)$ and cross-parcel capitalization $\phi_{ij}(z)$: the change in the land value of parcel j caused by what z

[4]Costs $c(z, P)$ are paid to outsiders (construction, financing spreads, governance, transactions) and are not part of the property's assessable value.

places on parcel i, with $\phi_{ii} \equiv 0$.

For each partition P, let $\sigma_b(z, P)$ denote the nonland market value assigned to block $b \in P$. This is an accounting allocation across assessment units, not a primitive assignment of each structure to a primitive parcel, and it satisfies $\sum_{b \in P} \sigma_b(z, P) = s(z)$. The market value of block b is

$$M_b(z, P) \equiv \sum_{i \in b} \ell_i + \sigma_b(z, P) + \sum_{i \in b} \sum_{j \neq i} \phi_{ji}(z).$$

The market value of the footprint is

$$M(z) \equiv \sum_{b \in P} M_b(z, P) = \bar{\ell} + s(z) + \sum_{i \neq j} \phi_{ij}(z),$$

which depends only on the physical development, not on assessment lines.

Assessment. Most parcels are built on, so land value must be imputed. Standard practice values a block's land as if vacant *with the surroundings as they stand*: the block's own improvements are removed in the counterfactual, while improvements on neighboring blocks remain standing. The assessed land value of block b is therefore

$$L_b(z, P) \equiv \sum_{i \in b} \ell_i + \sum_{i \in b} \sum_{j \notin b} \phi_{ji}(z),$$

and improvements are assessed residually, so that the two columns exhaust market value, $L_b(z, P) + I_b(z, P) = M_b(z, P)$. The land base is $R(z, P) \equiv \sum_{b \in P} L_b(z, P)$.

Proposition 1 (Assessment accounting)**.** *For any development z and partition P,*

$$R(z, P) = \bar{\ell} + \sum_{\substack{i \neq j:\ i,j \text{ in} \\ \text{different blocks}}} \phi_{ij}(z), \qquad M(z) - R(z, P) = s(z) + \sum_{\substack{i \neq j:\ i,j \text{ in} \\ \text{the same block}}} \phi_{ij}(z).$$

Total assessed value is partition-invariant. The partition determines only which capitalization terms are booked to the land column rather than the improvement column: a spillover enters the land base if and only if it crosses an assessment boundary.

Taxes. A recurring tax on assessed land and improvement values capitalizes into liabilities with present value Λ_L and Λ_I per dollar of assessed value, $0 \leq \Lambda_I \leq \Lambda_L \leq 1$.[5] The

[5]An annual rate τ at discount rate r capitalizes to $\Lambda = \tau/(r + \tau)$ when the base tracks market prices, which the tax itself depresses; the Georgist tax of 100% of land rent corresponds to $\Lambda_L = 1$. Practical proposals are not small in these units: a 5% annual land tax at $r = 5\%$ gives $\Lambda_L = 0.5$.

footprint's total liability is

$$T(z,P) = \Lambda_I M(z) + (\Lambda_L - \Lambda_I)\, R(z,P).$$

A pure LVT has $\Lambda_I = 0$; an ordinary property tax has $\Lambda_L = \Lambda_I$; a split-rate tax lies in between.

Corollary 1 (Boundary wedge)**.** *The partition affects the tax bill only through* $(\Lambda_L - \Lambda_I)\, R(z,P)$*. An ordinary property tax is therefore boundary-blind, and every boundary incentive below scales with the rate differential* $\Lambda_L - \Lambda_I$*.*

Decision problem. A decision-maker who internalizes the whole footprint, whether one owner, commonly controlled entities, or a Coasean coalition with side transfers (Appendix B), chooses the partition and a feasible development to maximize after-tax wealth,

$$W(z,P) \equiv M(z) - c(z,P) - T(z,P), \qquad z \in \mathcal{Z}(P).$$

Substituting Proposition 1 into $T(z,P)$ writes the objective component by component,

$$W(z,P) = (1-\Lambda_L)\left[\bar{\ell} + \sum_{\substack{i\neq j:\ i,j \text{ in} \\ \text{different blocks}}} \phi_{ij}(z)\right] + (1-\Lambda_I)\left[s(z) + \sum_{\substack{i\neq j:\ i,j \text{ in} \\ \text{the same block}}} \phi_{ij}(z)\right] - c(z,P)$$

so the owner retains each component of value at one minus the rate of the column where it is booked. In particular, a spillover keeps $(1-\Lambda_I)$ of its value while internal to a block and $(1-\Lambda_L)$ once it crosses a boundary. With no tax, the chosen configuration maximizes pre-tax value $M(z) - c(z,P)$; call such configurations *efficient*.[6]

3 The Partition Margin

Fix the development z. By Proposition 1, wealth varies with the partition only through the cross-boundary capitalization terms, so tax minimization sorts boundaries by the sign of spillovers.

[6]Efficiency is relative to the footprint, since spillovers within N are internalized by the decision-maker. An ordinary property tax is boundary-blind by the corollary but still discourages structures through $\Lambda_I M(z)$, though that familiar margin is not our subject.

Proposition 2 (Partition margin)**.** *Fix z, with boundaries technologically neutral. Merging blocks removes their mutual capitalization from the land base; splitting a block inserts the newly separated terms. If all $\phi_{ij}(z) \geq 0$, the land base is minimized by the single-block partition, where $R = \bar{\ell}$; if all $\phi_{ij}(z) \leq 0$, it is minimized by the finest partition. An LVT thus favors assembly under positive spillovers and fragmentation under negative ones.*

Corollary 2 (Pure reclassification)**.** *Under technologically neutral boundaries, a repartition $P \to P'$ leaves the development, market value, and total assessed value unchanged and changes wealth by $-(\Lambda_L - \Lambda_I)\big[R(z, P') - R(z, P)\big]$. Boundary changes are then pure tax arbitrage.*

When boundaries are technologically neutral, this margin misallocates nothing by itself. The same structures stand either way, and only paperwork changes. (The tax can still distort *which* structures stand, the development margin of Section 4.) Its costs are fiscal.

Corollary 3 (Revenue erosion)**.** *With costless repartition by decision-makers who internalize liabilities, every positive capitalization term can be merged out of the land base and every negative term deducted from it, so equilibrium revenue is at most $\Lambda_L \bar{\ell}$.*

The Henry George revenue claim therefore fails through boundary reclassification alone, before any real behavior changes. The differential rents the theorem taxes at the efficient city size are created by proximity itself, the ϕ terms rather than $\bar{\ell}$, so what erodes is exactly what the theorem needs. This reclassification is also what assessment data can detect, and Section 5 tests precisely this response of the assessed land base to splits and mergers.

The partition margin acquires welfare content when partitions change what can be built. Let a fine partition P^F permit development z^F and a unified partition P^U permit development z^U. Suppose the fine-grain project is efficient: $\Delta V \equiv [M(z^F) - c(z^F, P^F)] - [M(z^U) - c(z^U, P^U)] > 0$. But if it creates cross-boundary positive capitalization $\Phi \equiv R(z^F, P^F) - R(z^U, P^U) > 0$, a pure LVT makes the after-tax gap $\Delta V - \Lambda_L \Phi$. Thus whenever $0 < \Delta V < \Lambda_L \Phi$, the owner builds the lower-value unified project.

Example. Two adjacent parcels carry combined baseline land $\bar{\ell} = 100$. Option F keeps them on separate titles, a day care beside the apartments it serves, a layout the day care's own financing and resale require; option U drops the day care for a single apartment block on one title. The day care and the apartments raise each other's land by a combined $\Phi = 20$, a fifth of the base, and F is worth $\Delta V = 10$ more before tax. Without the LVT the owner builds F. But F exposes the $\Phi = 20$ to the tax, so a pure LVT makes the gap $10 - 20\Lambda_L$, and once $\Lambda_L > \frac{1}{2}$, an annual land-tax rate above the discount rate, she abandons the day care for the inferior block, sacrificing the positive spillover to escape the tax it would bear. Negative internal spillovers reverse the force, rewarding the fragmentation of holdings that should stay whole. Pre-fixing boundaries by regulation does not restore efficiency (Appendix A).

4 The Development Margin

The second margin is what gets built, and it costs welfare even when boundaries are technologically irrelevant. Consider a pure LVT with capitalized rate $\Lambda \equiv \Lambda_L$, $\Lambda_I = 0$, and write $X(z, P) \equiv R(z, P) - \bar{\ell}$ for the capitalization that development z exposes across the boundaries of P.

Proposition 3 (Spillover reweighting)**.** *Fix a partition P and $\Lambda \in [0, 1]$. The after-tax development choice solves*

$$\max_{z \in \mathcal{Z}(P)} \big[M(z) - c(z, P)\big] - \Lambda\, X(z, P)$$

and departs from the no-tax choice only through the exposure term $\Lambda X(z, P)$. For any $A, B \in \mathcal{Z}(P)$,

$$W(A, P) - W(B, P) = \big[M(A) - c(A, P)\big] - \big[M(B) - c(B, P)\big] - \Lambda\big[X(A, P) - X(B, P)\big],$$

so the LVT moves a development down the ranking in proportion to the positive capitalization it exposes and up in proportion to the negative capitalization it exposes. The ranking of A over B reverses exactly when the no-tax gap between them is positive but smaller than $\Lambda\big[X(A, P) - X(B, P)\big]$.

Proof. Setting $\Lambda_I = 0$ in the objective decomposition of Section 2 gives $W(z, P) = \big[M(z) - c(z, P)\big] - \Lambda X(z, P) - \Lambda\bar{\ell}$. The constant $\Lambda\bar{\ell}$ does not affect the maximizer; subtracting gives the pairwise expression. □

The effective price of spillovers. The same $\phi_{ij}(z)$ enters wealth twice, once through the asset value $M(z)$ and once, when exposed, through $X(z, P)$. The exposure is itself chosen. With boundaries technologically neutral the owner merges positive terms into blocks and splits negative ones out, so whenever some partition cuts every negative term without cutting a positive one, each spillover enters W at

$$\kappa(\phi) \equiv \begin{cases} \phi & \text{if } \phi \geq 0\text{, merged and untaxed,} \\ (1-\Lambda)\,\phi & \text{if } \phi < 0\text{, split and rebated,} \end{cases}$$

whereas the no-tax benchmark prices every term at ϕ. The owner partitions defensively against positive spillovers and opportunistically toward negative ones: merging keeps a gain untaxed, but splitting turns a loss into a rebate. Only the second is profitable, so the re-ranking of Proposition 3 runs one way, toward more damaging developments. If all spillovers are nonnegative, the single-block partition gives $X = 0$ for every development, and the LVT does not distort what gets built.

LVTs might still discourage positive spillover developments. Consolidation raises $c(z, P)$ when titles must be merged and held, and when fine boundaries are technologically necessary, as in Section 3, the positive-side penalty binds and cross-boundary amenities are taxed at rate Λ. The LVT thus mutes, at rate Λ, the price signal land markets attach to neighborhood effects, an inverse Pigouvian wedge with full discounting at the canonical 100% Georgist rate.

Example. Now sever the tie between boundaries and building from the previous example: let every development be feasible under every partition, so redrawing titles is pure paperwork. That freedom dissolves the positive-spillover distortion, since the day care and apartments can now share a single deed, internalizing the amenity untaxed while keeping the superior layout. A negative spillover is not cured, because there the owner gains from exposing the spillover rather than from hiding it. A lot sits beside those apartments, again with $\bar{\ell} = 100$. The owner can build a clean warehouse, or a gas station whose fuel tanks are privately better by $\Delta s = 10$ but lower the apartment land by $|\phi| = 20$, a fifth of the base. Internalizing that loss, she builds clean without the tax, since $10 < 20$. Under a pure LVT she splits the station onto its own title so the nuisance crosses the boundary, setting $X = -20$ and recovering 20Λ; the station then wins once $\Lambda > \frac{1}{2}$, again the point where the tax rate overtakes the discount rate. Boundaries change nothing real here, yet the tax still flips what gets built. Because she bears the loss in full but recovers only $\Lambda|\phi|$, poisoning her own land never pays on its own; the LVT subsidizes only nuisances whose private advantage already lies between $(1-\Lambda)|\phi|$ and $|\phi|$.

The distortion extends to victims of outsiders. A neighbor whose land value falls by $|\phi|$ recovers $\Lambda|\phi|$ in reduced taxes, so the most she will pay in bargaining, court, or zoning fights to stop a nuisance shrinks to $(1-\Lambda)|\phi|$; at $\Lambda = 1$ she is indifferent to the blight next door. Coasean discipline on neighborhood quality weakens in exact proportion to the tax (Appendix B).

5 Evidence from Assessment-Boundary Changes

An operational LVT taxes an assessed land component, not an observed bare land value. As discussed, capitalization of nearby amenities or nuisances into property values is a well-documented phenomenon. The empirical question is whether that capitalization is registered specifically in the assessor's reported land component when assessment boundaries shift.

In standard property-tax assessment practice, current market value is the benchmark for valuation (International Association of Assessing Officers, 2017, 2010). In Cook County, the Assessor states that land is generally valued per square foot at market value, and that property values are produced through mass appraisal using market sales, location, land and building characteristics, and related property data (Cook County Assessor's Office, 2026b,d,e). This is consistent with the capitalization mechanism in the model: if local amenities and nuisances affect transaction prices, a market-value assessment process may transmit those spillovers into assessed land values.

Documentation alone, however, does not let us infer whether actual assessments follow the counterfactual logic in the model. Land rates may be set at a neighborhood level, smoothing away parcel-specific spillovers; capitalization may be allocated to improvements rather than land; or reassessment lags may prevent short-run boundary changes from appearing in the reported land component. The empirical question is therefore not whether assessors are instructed to value land at market value. It is whether the implemented assessment process capitalizes local spillovers in a boundary-contingent way.

Concretely, the model predicts how the assessed land base should move when a boundary exposes or absorbs a local spillover. A split turns an internal interaction into a cross-boundary one, so the footprint's assessed land base should rise for positive-spillover pairs and fall for negative-spillover pairs. A merger does the reverse: it absorbs a cross-boundary interaction into a single assessment unit, so the assessed land base should fall for positive-spillover pairs and rise for negative-spillover pairs. These sign predictions on the assessed land base are the empirical targets.

Testing them matters even though no U.S. jurisdiction runs a pure LVT. The distortions identified in Sections 3–4 require that assessors price cross-boundary spillovers into land and net out within-boundary ones. If assessors do not do this, the distortion channel is shut down regardless of the tax rate. If they do, the question becomes one of magnitude rather than existence.

5.1 Data and Measurement

The empirical analysis uses publicly available annual Cook County, Illinois parcel maps and assessment rolls from 2001 through 2023 to identify 7,330 simple one-sided boundary changes: one parcel becoming two ($1 \to 2$) or two becoming one ($2 \to 1$). Values are summed over the full event footprint, so splits compare the parent to its children and mergers the parents to the merged child.

The analysis also uses four auxiliary sources: assessor characteristics and class files (proxies for pre-existing structures and pre-event development intensity), Chicago building permits (used to flag redevelopment and legal-reconfiguration events), environmental records (nuisance exposure from complaints, inspections, enforcement actions, storage tanks, and asbestos), and business-license and facility overlays (positive *service-care* uses, mostly daycare and preschool facilities and small clinics rather than large hospitals). Positive and nuisance tags use only pre-event activity, whereas redevelopment screens use a wider 5 year permit window.

To understand what drives the legal boundary changes and to identify confounders that could spuriously generate the predicted patterns, all 7,330 simple events were classified by a large language model.[7] Each event was passed to the model with its permit, environmental, business-license, and facility records, plus short summaries of the matched activity, and assigned a primary stated reason. Splits are mainly sale or ownership reorganizations (52.4%) plus administrative boundary cleanup (30.5%), followed by redevelopments (7.7%).[8] Mergers are even more dominated by routine boundary cleanup or class-change repartition (57.1%), followed by sale-linked reorganizations (32.5%). Appendix D.1 uses these codes as a robustness check, dropping events the model classifies as redevelopment.

To assess ownership, I use historical Cook County taxpayer names and mailing addresses as an ownership proxy. In the sample, 51.9% of mergers with complete pre-event

[7]The model used was GPT 5.4 with thinking effort set to `medium`.

[8]An administrative boundary cleanup is a paper reconfiguration in which the assessor reissues PINs to align tax-roll geometry with deeded lot lines, to correct prior mapping errors, or to follow a property-class reclassification, without any change in ownership, structure, or use. These are the cases where the model should apply most directly.

information already combine parent parcels held by the same owner. On the split side, 36.7% of events leave both child parcels under the same owner in the first post-change year, and 26.1% still do one year later. Countywide, 22.9% of 2025 parcels with usable ownership information belong to owner groups with at least two parcels on the same tax block.

The target outcome for the regressions is the county assessor's annual decomposition of assessed value into land and improvements, specifically log assessed land value per square foot. Assessed land values are noisy proxies for true bare-land values, but they are also the administratively relevant object, since any workable LVT would require assessors to impute bare land separately from structures.

The unit of valuation is a tax parcel, identified by a Property Index Number (PIN), and the Assessor defines a division as the legal or geographic reconfiguration of those parcel boundaries (Cook County Assessor's Office, 2026b). For each PIN and year, the Assessor reports separate land, building, and total assessed values (Cook County Assessor's Office Data Department, 2026). Cook County's roughly 1.8 million parcels are reassessed on a triennial cycle, with residential values produced by a statistical model and reviewed by analysts at the neighborhood level (Cook County Assessor's Office, 2026a,e,c).

Finally, to run the difference-in-differences design, I add stable never-treated controls to the event panel. Eligible controls are observed throughout the event window, never appear in any boundary event, retain the same assessment class throughout, and have no Cook County sale during the window. The panel caps each scope at 25 controls per treated event to keep weights roughly balanced across stacks. The widest sample draws controls countywide; tighter scopes restrict to the same township, the same assessor neighborhood, and, where support permits, the same neighborhood matched on broad class and area.

Further details on data and sample construction can be found in Appendix C.

5.2 Empirical Design

The baseline empirical design is a stacked difference-in-differences event study. Each stack contains one treated boundary-change footprint and its stable never-treated controls observed from $t-2$ through $t+2$, where t is the year of the split or merger event. The event-study specification estimates

$$y_{ise} = \alpha_i + \gamma_{\text{township}(i)\times s} + \sum_{\ell\in\{-2,0,1,2\}} \beta_\ell \mathbf{1}\{i \text{ is treated in stack } e\}\mathbf{1}\{s - t_e = \ell\} + \varepsilon_{ise},$$

where y_{ise} is log assessed land value per square foot, α_i is a site fixed effect, and $\gamma_{\text{township}(i)\times s}$ is a township-year time fixed effect. The year $t-1$ is omitted. Standard errors are clustered by boundary event. The pooled specification uses the same fixed effects but collapses the post period and compares it to the pre period.

Identifying assumptions. The main identifying assumption is parallel trends: absent the boundary change, treated footprints and their controls would have followed the same log-land-value path. A second assumption is no spillovers from the boundary change to the controls. This concern is relevant when a split or merger is bundled with redevelopment or another physical change on the treated footprint, since the associated amenity or nuisance shock may capitalize into nearby parcels used as controls. Depending on the spillover sign and on whether the event is a split or merger, spillovers to controls can attenuate or amplify the treated-control gap. To address this issue, the main nuisance sample screens out redevelopment-tagged events, and Appendix D.1 shows the estimates are robust to excluding redevelopment and related vertical-reconfiguration cases.

Measurement of spillovers. Spillovers are not directly observed, so the sign of the interaction between the parcels involved must be proxied. Nuisance tags use pre-event complaints, inspections, enforcement, storage-tank records (Chicago's registry of on-site fuel and chemical tanks tracked for soil-contamination risk), and asbestos-related records; positive tags use pre-event *service-care* activity, especially daycare, preschool, small clinics, and community-health or human-service facilities.

The main concern is that tags might be correlated with other relevant events. This is mitigated by the structure of the prediction itself, because the model predicts a four-cell pattern of signs rather than a one-way level effect: positive-spillover splits should raise assessed land values and positive-spillover mergers should lower them, while nuisances or negative externalities flip these signs. Generic confounders – a split shock, a merger shock, a tag-correlated neighborhood trend – cannot reproduce this product-sign pattern. Only an omitted factor that interacts jointly with the tag *and* with whether the event is a split or a merger would be a threat to identification.

Selection into reconfiguration. Splits and mergers are endogenous objects of optimization. Landowners reconfigure parcels when a different boundary structure raises the value of the underlying bundle, whether by changing the timing or scale of development, improving marketability, or enabling land assembly (Titman, 1985; Brooks and Lutz, 2016).

The target of this optimization is the joint value of the property, not the assessor's imputed land component, and the two need not move together (Clapp and Lindenthal, 2022). Selection is therefore a real concern, but it does not on its own generate the signed pattern the model predicts.

A generic optimization story implies that voluntary reconfigurations weakly raise property value net of reconfiguration costs; it says nothing about why the land component should fall in nuisance splits and service-care mergers but rise in the symmetric cells. To confound the sign tests, the selection process would have to interact jointly with spillover sign and with whether the event is a split or a merger. Appendix D.1 reports several exercises that address this residual concern.

5.3 Results

Figure 1 shows the headline signed event studies using the main control sample. The signing convention leaves split outcomes in their raw scale and multiplies merger outcomes by -1. As a result, the assessment model predicts negative coefficients for nuisance-tagged events and positive coefficients for positive-spillover events. After the boundary change, the nuisance sample falls from -0.148 in year t to -0.238 by $t+2$. The *service-care* sample jumps to 0.276 in year t, and remains positive and significant. The pooled all-stable DiD coefficients are -0.151 for nuisance events and 0.253 for *service-care* events.

Figure 1: Stacked DiD event studies for nuisance and positive-spillover events

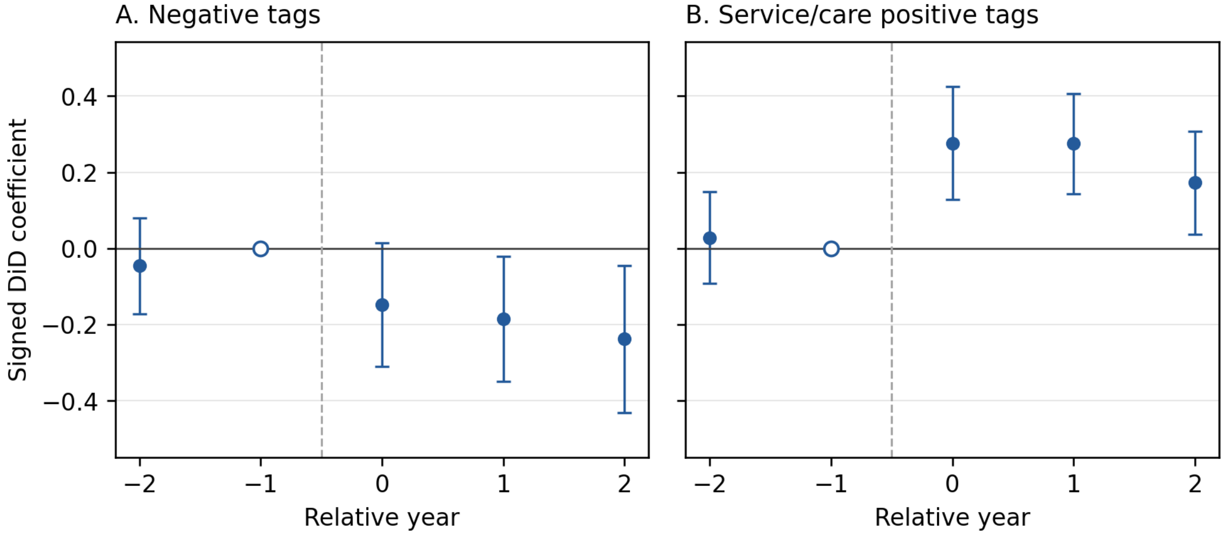

Notes: The outcome is log assessed land value per square foot. Coefficients are signed by event type: split outcomes are left in the raw scale and merger outcomes are multiplied by -1. Thus theory-consistent effects are negative for nuisance tags and positive for positive-spillover tags. Panel A uses nuisance-tagged Chicago events with all pre-event components in Chicago and no redevelopment or legal-reconfiguration screen hits ($N = 1{,}560$ event-year observations). Panel B ($N = 4{,}290$) shows results for *service-care* positive-use events. Both panels use all stable controls, site fixed effects, township-by-year fixed effects, and event-clustered standard errors. Whiskers are 95% confidence intervals.

Figure 2 reports the pooled post-pre DiD in two ways. The top row shows the raw split and merger coefficients, while the bottom row applies the same event-type signing used in Figure 1. The raw signs reverse as predicted by theory: in the nuisance sample, split coefficients are negative and merger coefficients are positive; in the *service-care* sample, split coefficients are positive and merger coefficients are negative. The signed pooled panels then place the nuisance effects below zero and the positive-spillover effects above zero. The same-township and same-neighborhood estimates are close to the all-stable estimates, which is reassuring because these local samples are designed to absorb common reassessment and neighborhood shocks.

The magnitudes are economically meaningful. Pooled effects of about -0.15 for nuisance events and 0.25 for positive-spillover events are large enough to shift parcel-configuration incentives under the high land-tax rates considered in practical LVT proposals. The assessed land base moves when assessment boundaries expose or absorb local spillovers, and it moves in the direction predicted by the assessment model.

Several robustness checks reported in Appendix D support this interpretation. The

Figure 2: Raw and signed pooled DiD effects

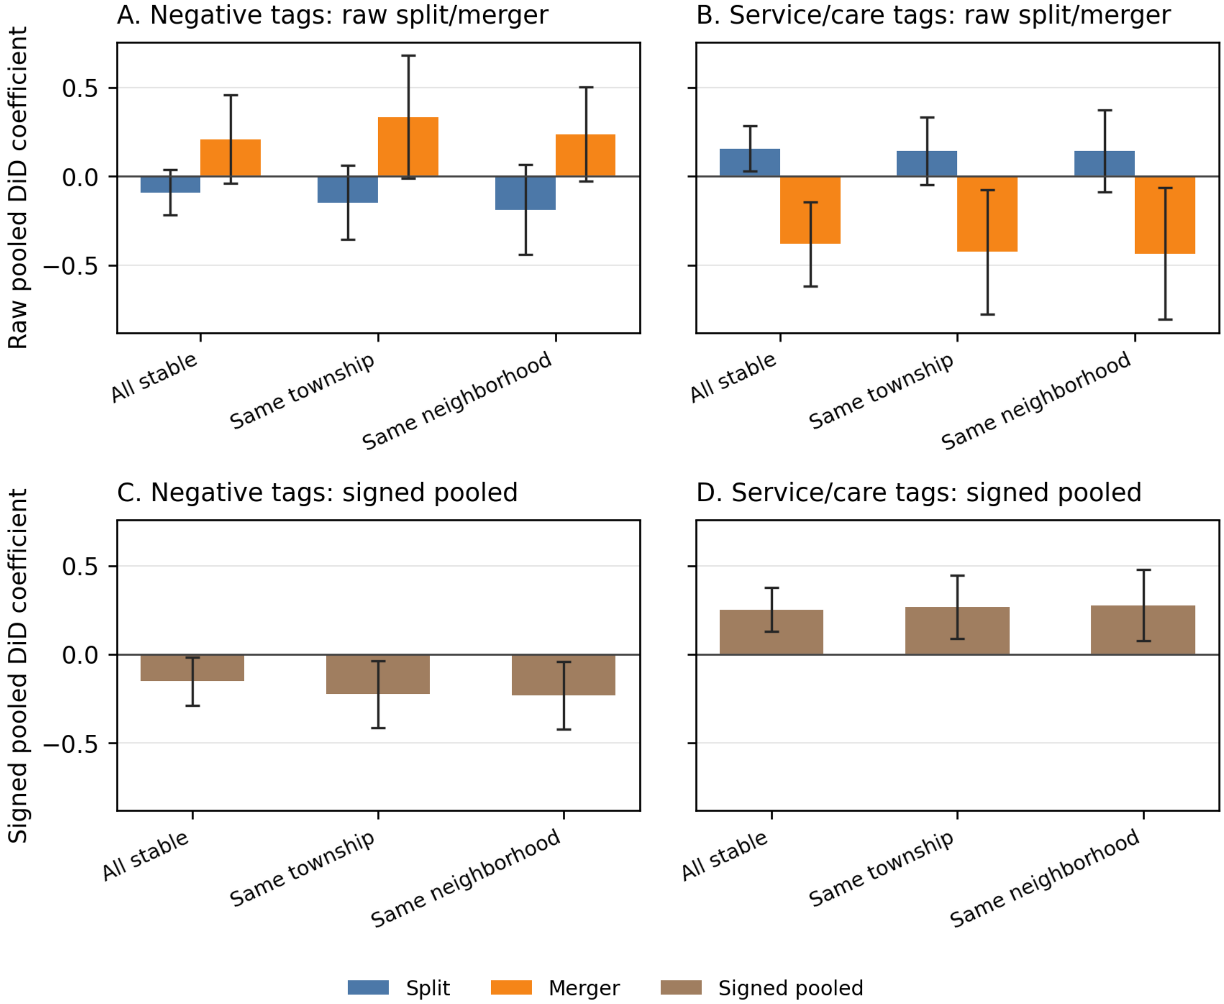

Notes: Bars report the pooled pre/post DiD coefficient comparing post years $t, t+1, t+2$ to pre years $t-2, t-1$. Panels A and B show raw split and merger estimates. Panels C and D sign the same pooled estimates by event type, leaving splits in raw scale and multiplying mergers by -1. Event-year observations: nuisance $1{,}560$ (all-stable and same-township) and $1{,}230$ (same-neighborhood); *service-care* $4{,}290$ (all-stable and same-township) and $3{,}535$ (same-neighborhood). The narrower neighborhood-class-area sample is omitted here because it has too little nuisance support. Whiskers are 95% confidence intervals.

four-sign pattern is concentrated in the land component rather than total or improvement assessed value; it is unchanged when redevelopment and condo-vertical events are dropped; and placebo tests on untagged neutral parcel reconfigurations do not reproduce it. The estimates are also robust to alternative DiD estimators (Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfoeuille, 2020; Sun and Abraham, 2021), to restricting controls to the same township or assessor neighborhood, and to broadening the positive-spillover tag beyond service-care.

6 Toward a Boundary-Robust LVT

By Proposition 1, the operational land base is a baseline $\bar{\ell}$ that no private choice can move plus cross-boundary capitalization terms that partitions and developments both move. The distortions of Sections 3–4 all run through the second piece, so a first-best correction must make the liability that owners can minimize independent of their choices.[9] Two information sets suffice.

The first is direct pricing of spillovers. If the planner observes every $\phi_{ij}(z)$, the land base can be reduced to the baseline $\bar{\ell}$ while each spillover is taxed or credited at source. The boundary arbitrage disappears because the base no longer depends on which terms cross boundaries, and the development distortion disappears because neighborhood effects face their social price rather than the subsidy of Section 4. This correction also aligns incentives for owners outside any coalition, since Pigouvian pricing does not depend on who owns which block.

The second is owner- or coalition-invariant assessment. If the planner observes which parcels share a decision-maker, whether through ownership, anticipated bargaining, or post-split coordination, those sets can be assessed as single fiscal units. Within-coalition ϕ_{ij} terms then drop out of the base however titles are drawn, removing at once the boundary arbitrage, the distortion of ownership structure, and the rebate on self-inflicted damage. What survives is the wedge between coalitions. Victims of an outside nuisance still recover share Λ_L of their loss, so private resistance to nuisances stays weakened across coalition lines even as it is restored within them.

[9] This invariance is also necessary: any corrected scheme whose minimized liability still varies with the development can reverse the development ranking for some configuration of private values.

7 Conclusion

This paper identifies a concrete case where land value taxes lose their canonical efficiency. The mechanism is straightforward. Parcel improvements generate spillovers that affect neighbors' assessed land values, and these spillovers enter the land tax base only when they cross assessment boundaries. Consolidating parcels internalizes them and removes them from the tax base, so the base itself depends on how land is partitioned.

This results in two distortions. First, owners can lower their tax bill by merging parcels when spillovers between them are positive and splitting when they are negative. Second, owners can split parcels to expose negative spillovers across an assessment boundary, lowering the assessed land base. The tax then refunds a share of the damage equal to the capitalized rate, so developments whose private advantage falls just short of the harm they inflict are tipped from rejected to built. Both require that one decision-maker internalize tax liabilities across multiple units: through common ownership, Coasean bargaining, or a split that creates the cross-boundary exposure in the first place.

The paper also tests the model's key prediction about how assessments respond to boundary changes. Using 7,330 simple splits and mergers in Cook County, Illinois between 2001 and 2023, stacked difference-in-differences estimates show that assessed land values move in opposite directions for splits and mergers, and that the direction reverses between nuisance and positive-spillover exposures – exactly the pattern the model predicts. Cook County is not an LVT jurisdiction, but any operational LVT would have to impute bare-land values from the same kind of data. The evidence therefore speaks directly to the measurement environment an LVT would have to operate in.

These results do not overturn the general case for land value taxation, nor do they speak to the LVT's efficiency relative to other taxes. Where land values derive from broad locational advantages or public infrastructure, traditional efficiency arguments remain sound. But within these limits, the concern raised here is real: LVTs subsidize externalities that reduce the assessed land base, and the subsidy mutes the owner's private value loss when the boundaries that determine the tax base are themselves choice variables. First-best corrections exist, such as directly pricing spillovers or assessing commonly-controlled parcels as a single unit, but each requires information not freely available. Designing a boundary-robust LVT is therefore a mechanism-design problem, not just an assessment problem.

A Fixed Boundaries and Development Choice

Section 3 let owners draw their own boundaries. This appendix asks whether a regulator can restore efficiency by fixing assessment boundaries *ex ante*.

Regulator-fixed boundaries and the implementation problem. Fix a partition P and consider a pure LVT with capitalized rate Λ. Footprint wealth is

$$W(z,P) = M(z) - c(z,P) - \Lambda\big[\bar{\ell} + X(z,P)\big],$$

where $X(z,P)$ is the exposed capitalization defined in Section 4. The baseline term is constant, but the exposed capitalization still depends on the development. Unless those terms are constant in z, the induced development

$$z(P;\Lambda) \in \arg\max_{z\in\mathcal{Z}(P)} W(z,P)$$

differs from the no-tax choice $z(P;0)$: fixing boundaries stops the arbitrage over boundaries but not the distortion in what gets built. A regulator who anticipates this response faces a second-best problem,

$$P^{SB}(\Lambda) \in \arg\max_{P}\ M\big(z(P;\Lambda)\big) - c\big(z(P;\Lambda),P\big),$$

which need not coincide with the partition that is efficient under no tax. When boundaries are technologically relevant, the trade-off sharpens, because the partition that fits the technology may be exactly the one that exposes negative capitalization to the tax.

Benchmark: separate owners, fixed boundaries, no transfers. The fixed-boundary distortion above requires one decision-maker to internalize liabilities across multiple blocks. Suppose instead that each block $b \in P$ is separately owned and transfers are unavailable. Owner b's assessed land value,

$$L_b(z,P) = \sum_{i\in b} \ell_i + \sum_{i\in b}\sum_{j\notin b} \phi_{ji}(z),$$

does not depend on her own building choices: her improvements are removed in her block's vacancy counterfactual, and their value sits in the untaxed improvement column. She therefore builds exactly what she would build with no tax. With fixed boundaries, separate owners, and no transfers, a pure LVT does not distort within-block development choice; the distortions in the main text require one decision-maker to internalize tax liabilities across blocks, through common ownership or bargaining with transfers.

B Coase Benchmark

The decentralized benchmark in Appendix A relied on the absence of transfers: each owner chose her own block's development and ignored the tax consequences imposed on neighbors. This appendix considers the opposite extreme, in which owners can bargain costlessly over the partition and the development and can make lump-sum side transfers.

Proposition 4 (Coase benchmark)**.** *Suppose owners can costlessly bargain over (z, P) and make lump-sum side transfers. Then, regardless of the initial ownership assignment, voluntary bargaining implements some*

$$(z^C, P^C) \in \arg \max_{P,\, z \in \mathcal{Z}(P)} W(z, P),$$

the single-decision-maker outcome of the main text.

Proof. Take any status quo. A proposed move to (z, P) changes the coalition's total payoff by $W(z, P)$ minus the status-quo value, and side transfers are zero-sum. If the gain is positive, a transfer scheme makes every owner weakly better off and at least one strictly better off; if it is nonpositive, no deviation is unanimously accepted. The stable bargaining outcomes are therefore exactly the maximizers of W. $\square$

Bargaining thus reproduces the boundary arbitrage and the reweighted development choice of the main text without common ownership; ownership affects only how the payoff is divided, not what is chosen.

Bargaining over nuisances. The tax also reshapes bargaining when no coalition forms. Consider a victim whose land value falls by $|\phi|$ if an outside owner builds a nuisance. Under a pure LVT at capitalized rate Λ, her after-tax loss is $(1 - \Lambda)|\phi|$, because the assessor deducts the damage from her land value and the tax refunds share Λ of it. Her maximum willingness to pay to stop the nuisance, whether in side payments, litigation, or political opposition, shrinks from $|\phi|$ to $(1 - \Lambda)|\phi|$, while the builder's reservation price is unchanged. Bargains that would have stopped the nuisance at any price in $\big((1-\Lambda)|\phi|,\ |\phi|\big)$

no longer occur, and at the Georgist limit $\Lambda = 1$ the victim is indifferent, since the tax authority absorbs the entire loss. Coasean enforcement of neighborhood quality therefore weakens in exact proportion to the tax, as stated in Section 4.

C Data Construction and Descriptive Statistics

All raw data underlying this paper come from publicly available Illinois open-data records. Cook County parcel maps and the historical parcel-history layer were downloaded from the county's public ArcGIS REST service. Cook County assessor characteristics, class files, parcel-address taxpayer names, and assessment rolls come from the Cook County open data portal. The City of Chicago open data portal supplies building permits, environmental records (complaints, inspections, enforcement, storage tanks, asbestos), business licenses, and facility overlays. Both portals serve their datasets through the Socrata API, so the project's collection scripts hit each endpoint directly and store raw extracts under the project's data tree. Each record is timestamped and the API returns the full historical row for every dataset, so the panel can be rebuilt from scratch.[10]

The main analysis sample is the annual event-level site panel built from Cook County parcel-history maps, with all assessed values summed over the full physical footprint of the event. Auxiliary nuisance and positive-use tags are assigned to pre-event site geometry, while redevelopment screens deliberately use a wider around-event permit window. Descriptive taxpayer-name ownership proxies for the pre-merge parent parcels and post-split child parcels draw from the Cook County parcel-address file.

D Additional Results and Robustness

This appendix collects four supporting exercises. The first addresses the remaining concern, raised in the main text, that the spillover tags and selection into reconfiguration could be confounded. The other three check robustness to alternative DiD estimators, to local control samples, and to broader positive-use tags.

[10] Cook County parcel-history GIS service: `https://gis.cookcountyil.gov/traditional/rest/services/parcelHistorical/MapServer`. Cook County data catalog: `https://datacatalog.cookcountyil.gov`. City of Chicago data portal: `https://data.cityofchicago.org`. Dataset identifiers, vintages, and the timing rules used to assign exposure to events are listed in Table 1.

Table 1: Empirical data sources and timing assignment

Source	Scope and years	Role in the analysis	Timing rule
Cook County parcel history and annual parcel maps	Countywide, 2000–2023	PIN births and exits, site footprints, summed assessed land and total value	$t-1$ polygons for pre-event site construction; t is first post-change parcel year
Cook County parcel addresses and taxpayer names	Countywide, 1999–2026	Common-ownership descriptives for split and merge events	Matched to pre-event parent parcels in $t-1$ and split child parcels in t and $t+1$
Cook County assessor characteristics and class files	Countywide, pre-event years	Multiple-buildings proxy, building intensity, residential-use profiles	Matched to pre-event PINs
Chicago building permits	Chicago, 2006–2023	Redevelopment screens and positive-use permit text	Positive-use text uses only pre-period permits in $[t-2, t-1]$; redevelopment screens use $[t-2, t+2]$
Chicago environmental records	Chicago, 2006–2023	Nuisance tags such as complaints, inspections, enforcement, tanks, and asbestos records	Pre-event exposure if active or recorded in $[t-2, t-1]$
Chicago business licenses	Chicago, 2002–2026	Positive-use tags such as daycare, preschool, clinics, and cafes	Active during $[t-2, t-1]$ and spatially matched to pre-event site geometry
Chicago facility overlays	Chicago, current layers	Breadth check for schools, libraries, parks, and care facilities	Used as reviewed enrichment, not as the benchmark positive-use design

Notes: Throughout the empirical section, t denotes the first post-change parcel year. Pre-event nuisance and positive-use exposures are assigned from activity already present on the year-($t-1$) parcel map during tax years $t-2$ and $t-1$. Assessed values are always summed over the full event footprint, so splits compare the parent site to the sum of the children and mergers compare the sum of the parents to the merged child.

Table 2: How common are boundary events, and why do they happen?

	Splits ($1 \rightarrow 2$)	Mergers ($2 \rightarrow 1$)
Total simple events, 2001–2023	4,492	2,838
Boundary or administrative cleanup	30.5	57.1
Sale or ownership reorganization	52.4	32.5
Redevelopment or new construction	7.7	3.1
Nuisance or environmental story	5.2	1.9
Positive-use or service amenity	1.3	0.8
Condo / vertical internal reconfiguration	0.4	0.8
Public or exempt institutional	1.4	0.5
Unclear other	1.1	3.2

Notes: The first row reports the simple-event universe used in the baseline empirical design. All remaining rows report percentages from the full LLM audit of all 7,330 simple events. Mergers are dominated by routine parcel housekeeping and class-change repartition, while splits have a larger sale-linked and redevelopment tail. This is why the paper leans on tagged subsamples rather than generic average-event comparisons.

Table 3: Common ownership around assessment-boundary changes

Statistic	Count / denominator	Share
Panel A. Countywide ownership proxy, 2025		
Parcels with usable ownership proxy	1,861,820 / 1,862,805	99.9%
Parcels in multi-parcel owner groups	489,411 / 1,861,820	26.3%
Parcels in owner groups with at least one same-tax-block pair	425,546 / 1,861,820	22.9%
Panel B. Event-level same-owner shares		
Mergers with same owner across parent parcels in $t-1$	1,410 / 2,718	51.9%
Splits with same owner across child parcels in t	1,488 / 4,050	36.7%
Splits with same owner across child parcels in $t+1$	1,019 / 3,904	26.1%

Notes: Panel A uses the 2025 Cook County parcel-address file and treats normalized taxpayer name plus mailing address as the ownership proxy when available, then owner name plus owner address, and finally the normalized name alone as a fallback. The same-tax-block measure uses the Cook County parcel-number block code (the shared `PINA-PINSA-PINB` prefix of the parcel number), so it captures owner groups with at least two parcels on the same tax block rather than owner groups that are merely dispersed countywide. Panel B applies the same ownership proxy to the boundary-event sample. Here t is the first post-change parcel year, so mergers compare the parent parcels in $t-1$, while splits compare the child parcels in t and again in $t+1$. Shares use only events with complete owner information on the relevant side. These ownership proxies are descriptive and conservative: exact-name matching alone delivers even higher same-owner rates.

D.1 Robustness to selection

Land, total, and improvement decomposition. The headline outcome is log assessed land per square foot, but the model's prediction is specifically about the land component because the LVT base is land-only. The most credible alternative story is that owners reconfigure to optimize the property bundle, and the assessor's land and improvement decomposition then reflects whatever shock made the reconfiguration valuable. Under that story, total assessed value and improvement value should track land in the same direction, since the assessor is slicing a single rising or even a falling pie. The mechanism in this paper, by contrast, predicts that spillover value between neighboring parcels gets reclassified between the land and improvement lines without changing the underlying property: the land component should move in the four-sign pattern, while total assessed value tracks the optimization story and improvement value moves in the opposite direction from land in the cells where the assessment boundary matters.

Figure 3 and Table 4 report the same pooled DiD for all three outcomes. In the two cells where the model predicts negative land effects – nuisance splits and service-care mergers – the land component is the most negative outcome by a wide margin: the service-care merger coefficient is -0.35 on log land per square foot but -0.12 on log total assessed value (insignificant) and $+0.11$ on log improvement value, and the nuisance split coefficient is -0.09 on log land per square foot, $+0.11$ on log total, and $+0.18$ on log improvement. In the two "soft" cells – nuisance mergers and service-care splits – the boundary mechanism and voluntary optimization push in the same direction, and all three outcomes move together in the theory-consistent direction.

This decoupling is the point of the test. A confounder that explained the headline land result without invoking the boundary mechanism would have to drive total and improvement value in the same direction as land within each cell; what we see instead is that the four-sign pattern is sharply concentrated in the land outcome and does not propagate to the rest of the property's assessed value.

Figure 3: Land, total, and improvement assessed-value response by cell

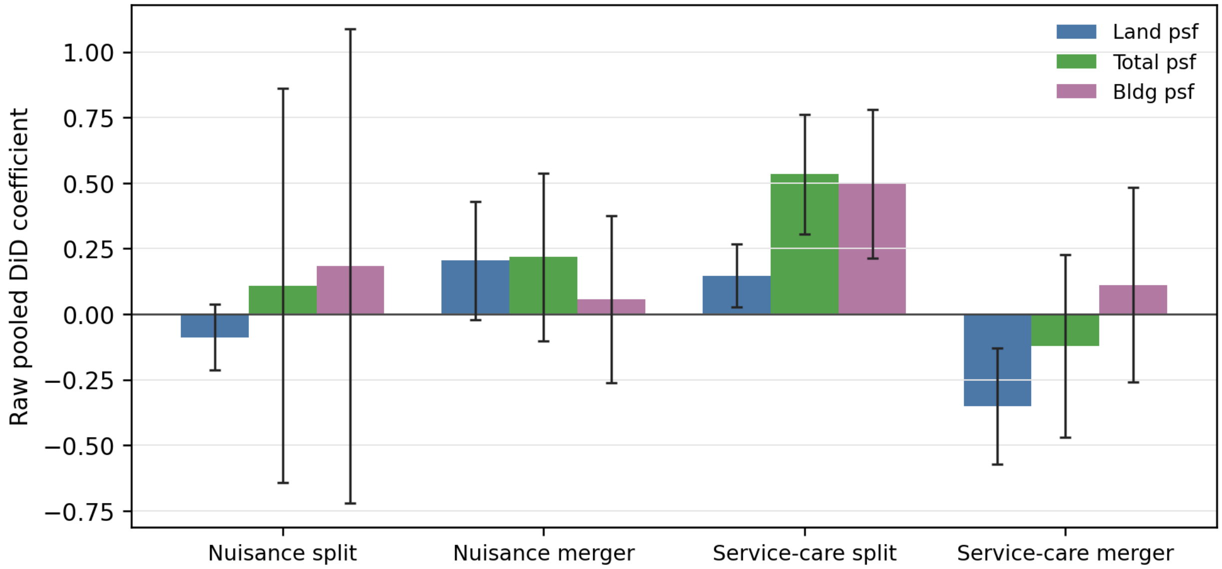

Notes: Pooled raw DiD coefficients on log assessed land per square foot, log total assessed value per square foot, and log improvement (building) value per square foot, by cell, all-stable control sample. Whiskers are 95% confidence intervals. Cell event-year observations: nuisance split $N = 1{,}040$, nuisance merger $N = 520$, service-care split $N = 2{,}730$, service-care merger $N = 1{,}560$. Site and township-by-year fixed effects are included throughout.

Table 4: Pooled DiD coefficients across land, total, and improvement value, by cell

Cell	N obs	Land psf	Total psf	Improvement psf
Nuisance split	1,040	-0.088 (0.064)	$+0.109$ (0.384)	$+0.184$ (0.462)
Nuisance merger	520	$+0.205$ (0.115)	$+0.218$ (0.164)	$+0.057$ (0.162)
Service-care split	2,730	$+0.147$ (0.062)	$+0.534$ (0.117)	$+0.497$ (0.144)
Service-care merger	1,560	-0.351 (0.112)	-0.121 (0.178)	$+0.111$ (0.189)

Notes: Raw pooled DiD coefficients with event-clustered standard errors in parentheses, all-stable control sample. Each cell is one tag-by-direction subset of the headline samples: nuisance no-redevelopment events split into splits and mergers, and service-care enriched-core events split into splits and mergers. Outcomes are log assessed land per square foot, log total assessed value per square foot, and log improvement (building) value per square foot. Site and township-by-year fixed effects are included for every cell-outcome cell. The two "hard" cells are nuisance splits and service-care mergers, where the model predicts negative land effects from boundary capitalization.

LLM-coded event-reason stratification. A separate full LLM audit independently classified all 7,330 simple boundary events by stated reason, without using the headline tag. All nuisance-tagged events fall in the audit's nuisance and environmental bucket; zero are independently classified as redevelopment or condo reconfiguration. The headline service-care sample is also dominated by the audit's positive-use bucket (28 of 33), with a small redevelopment tail (4 of 33) and one nuisance-classified event.

The most direct response to the concern that redevelopment drives the results is therefore to drop the events independently classified as redevelopment. Figure 4 and Table 5 report the result. The nuisance theory-signed coefficient is unchanged (-0.15 on log land per square foot in both the headline and drop-redevelopment samples, since no nuisance events are LLM-coded as redevelopment). The service-care coefficient barely moves: $+0.25$ in the headline becomes $+0.24$ after dropping the four LLM-coded redevelopment events, and the four redevelopment service-care events on their own carry a $+0.32$ coefficient, going in the same direction as the headline. The headline pattern is therefore not a reflection of redevelopment selection.

Figure 4: Theory-signed pooled DiD by LLM event-reason stratum

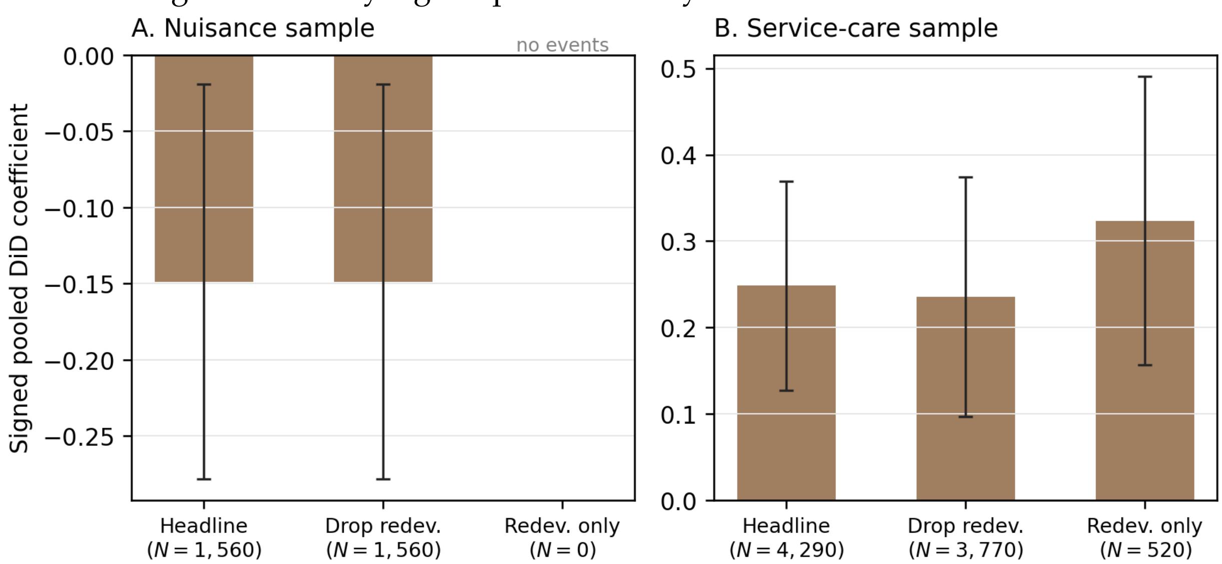

Notes: Theory-signed pooled DiD coefficients on log assessed land per square foot, all-stable control sample. Coefficients use the event-type signing in the main text, so theory-consistent nuisance effects are negative and theory-consistent service-care effects are positive. Panel A shows the nuisance no-redevelopment sample under three reason strata; the redevelopment-only stratum is empty because no nuisance events are LLM-coded as redevelopment or condo reconfiguration. Panel B shows the service-care enriched-core sample under the same three strata. Stratum event-year observations: nuisance headline 1,560, drop redevelopment 1,560, redevelopment only 0; service-care headline 4,290, drop redevelopment 3,770, redevelopment only 520. Whiskers are 95% confidence intervals.

Table 5: Theory-signed pooled DiD by LLM event-reason stratum

Sample	Reason stratum	N obs	Coef. [95% CI]
	Headline	1,560	-0.149 [-0.278, -0.019]
Nuisance	Drop redevelopment	1,560	-0.149 [-0.278, -0.019]
	Redevelopment only	0	—
	Headline	4,290	$+0.248$ [$+0.127$, $+0.369$]
Service-care	Drop redevelopment	3,770	$+0.236$ [$+0.097$, $+0.374$]
	Redevelopment only	520	$+0.324$ [$+0.157$, $+0.491$]

Notes: Theory-signed pooled DiD coefficients on log assessed land per square foot, all-stable control sample. Nuisance coefficients use the event-type signing convention (splits in raw scale, mergers multiplied by -1), so theory-consistent values are negative. The redevelopment stratum groups events the LLM independently classifies as either new-construction redevelopment or condo vertical reconfiguration; the "drop redevelopment" rows exclude these events from the headline sample, and the "redevelopment only" rows restrict to them. No nuisance events fall in the redevelopment stratum, so the headline and drop-redevelopment rows coincide for that sample. Confidence intervals use event-clustered standard errors. Site and township-by-year fixed effects are included throughout.

Placebo using untagged events. The third exercise repeats the same DiD on a placebo sample of *neutral* reconfiguration events: Chicago boundary changes that the data and audit do not flag as either nuisance or positive-use events, and that show no permit activity consistent with redevelopment, condominium conversion, or a purely legal reconfiguration of the parcel. From the events that meet these criteria, we draw a reproducible random subsample, seeded so the draw is the same each time the panel is rebuilt. Figure 5 and Table 6 report the raw split and merger coefficients alongside the tagged samples.

Neutral mergers move zero on average (-0.005, SE 0.034). Neutral splits show a small generic increase ($+0.071$, SE 0.028), which is consistent with a boundary-events-as-optimization story and on its own raises the question of whether the headline split coefficients partly reflect a generic per-square-foot accounting effect of subdivision. They are not, however, driven entirely by it: the nuisance split coefficient (-0.09) is still significantly below the neutral split baseline, and the service-care split coefficient ($+0.15$) is still above. The four-sign pattern in tagged events thus survives net of the generic split-direction baseline, and the merger side has no comparable baseline to net out.

Figure 5: Raw split and merger pooled DiD effects by tag group (placebo)

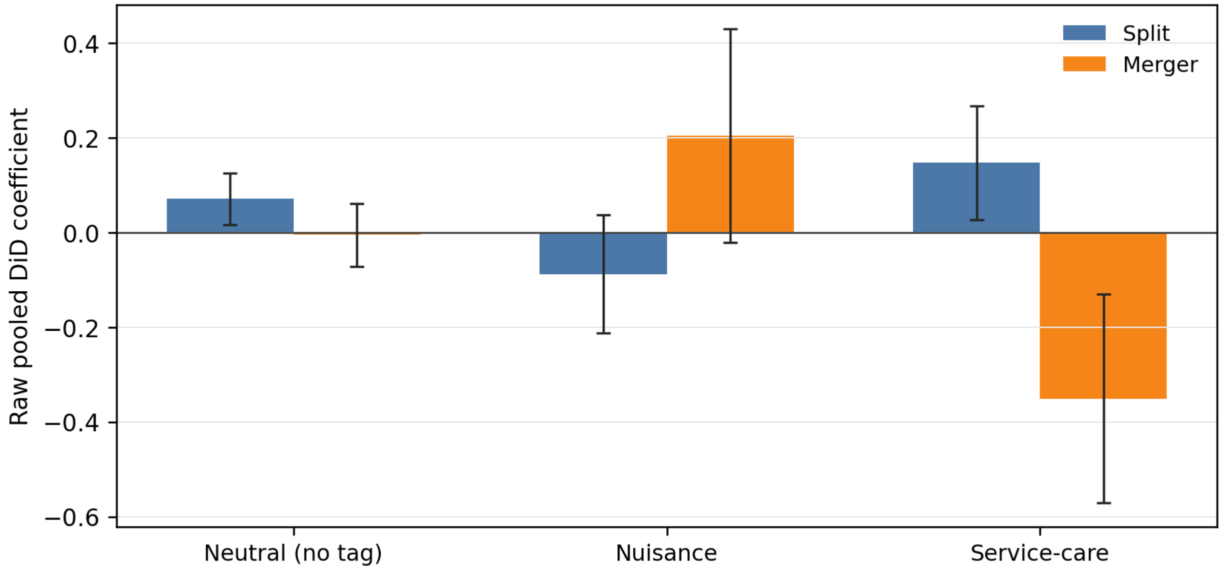

Notes: Pooled raw DiD coefficients on log assessed land per square foot, all-stable control sample. Bars are split (left) and merger (right) coefficients within each group. "Neutral" events are simple boundary changes with no nuisance evidence (no environmental complaints, inspections, enforcement, tanks, or asbestos records), no service-amenity evidence in the permit, business-license, or facility data, no redevelopment permit activity, and no flag for condominium conversion or a purely legal reconfiguration of the parcel. The sample is restricted to Chicago townships and event years ≥ 2006 to match the tagged samples, and is a reproducible random subsample drawn from the events meeting these criteria. Event-year observations: neutral $18{,}200$ (splits), $7{,}800$ (mergers); nuisance $1{,}040$ (splits), 520 (mergers); service-care $2{,}730$ (splits), $1{,}560$ (mergers). Whiskers are 95% confidence intervals.

Table 6: Raw split and merger DiD effects by tag group (placebo)

	Splits		Mergers	
Group	N obs	Coef. [95% CI]	N obs	Coef. [95% CI]
Neutral (no tag)	18,200	$+0.071$ $[+0.016,\ +0.125]$	7,800	-0.005 $[-0.072,\ +0.061]$
Nuisance	1,040	-0.088 $[-0.213,\ +0.037]$	520	$+0.205$ $[-0.021,\ +0.430]$
Service-care	2,730	$+0.147$ $[+0.027,\ +0.268]$	1,560	-0.351 $[-0.571,\ -0.131]$

Notes: Pooled raw DiD coefficients on log assessed land per square foot, all-stable control sample, $\text{post} = \{t, t+1, t+2\}$ vs. $\text{pre} = \{t-2, t-1\}$. "Neutral" events are simple boundary changes with no nuisance evidence (no environmental complaints, inspections, enforcement, tanks, or asbestos records), no service-amenity evidence in the permit, business-license, or facility data, no redevelopment permit activity, and no flag for condominium conversion or a purely legal reconfiguration of the parcel. The sample is restricted to Chicago townships and event years ≥ 2006 to match the tagged samples, and is a reproducible random subsample drawn from the events meeting these criteria. Confidence intervals use event-clustered standard errors. Site and township-by-year fixed effects are included throughout.

D.2 Alternative DiD estimators

Figure 6 overlays the baseline TWFE event-study coefficients with three alternative estimators designed for staggered adoption: Callaway and Sant'Anna (2021), de Chaisemartin and D'Haultfoeuille (2020), and the interaction-weighted estimator of Sun and Abraham (2021) using never-treated controls. The alternatives are not identical specifications: TWFE and Sun–Abraham include site and township-by-year fixed effects, while the other two implement their own never-treated-control comparisons on the same stacked sample. The goal is estimator robustness, not exact equality of specifications. The post-event pattern remains negative for nuisance tags and positive for *service-care* tags across estimators.

Figure 6: Alternative DiD estimators for the headline event studies

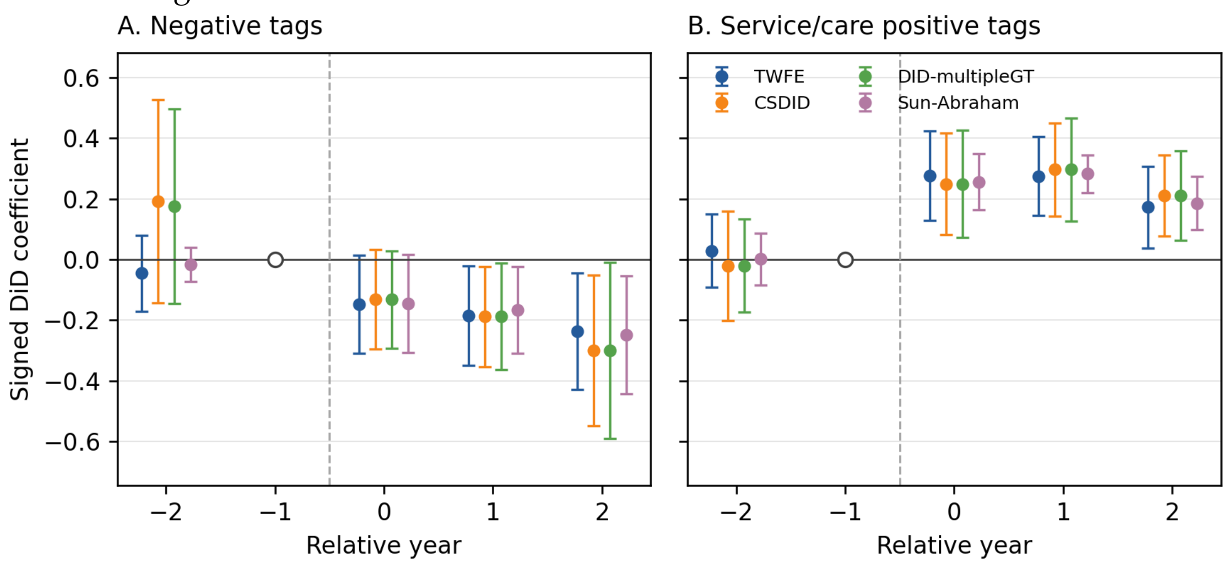

Notes: The figure repeats the all-stable event studies from the main text and overlays TWFE, Callaway–Sant'Anna, de Chaisemartin–D'Haultfœuille, and Sun–Abraham estimates. All estimators use the same stacked treated/control sample and never-treated stable controls. Coefficients use the event-type signing in the main text; the omitted $t-1$ period is shown as an open dot at zero. Panel A has $N = 1{,}560$ event-year observations (nuisance). Panel B has $N = 4{,}290$ (enriched-core *service-care*). Whiskers are 95% confidence intervals.

D.3 Local control samples

Figure 7 repeats the baseline TWFE event studies using local control samples. The post-event coefficients retain the predicted negative sign for nuisance tags and positive sign for *service-care* tags across both local scopes, which helps rule out broad township or neighborhood reassessment shocks as the main driver.

Figure 7: Event-study estimates with local stable controls

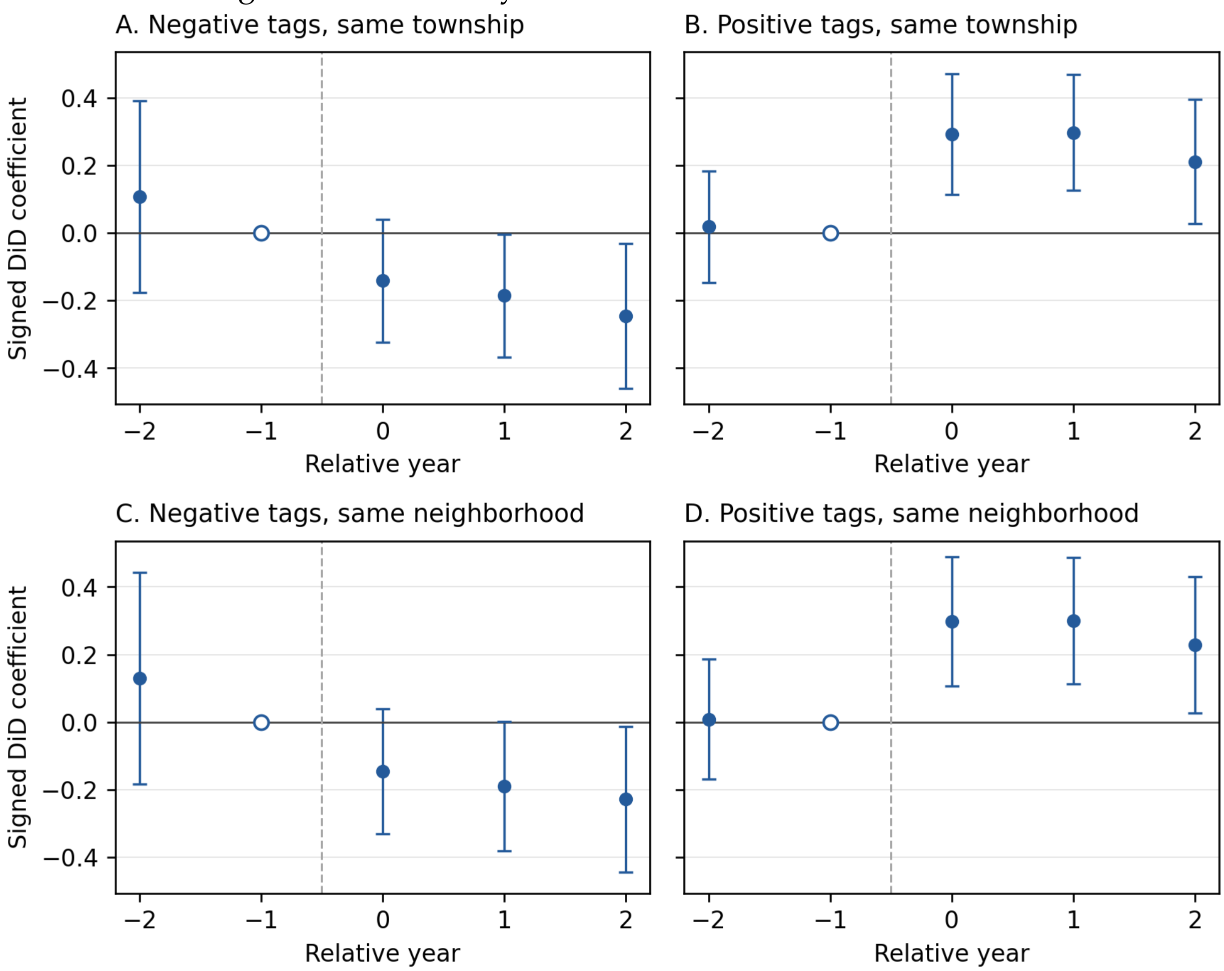

Notes: The figure repeats the baseline TWFE event studies using stable controls from the same township or same assessor neighborhood. Event-year observations: nuisance 1,560 (same-township) and 1,230 (same-neighborhood); enriched-core *service-care* 4,290 (same-township) and 3,535 (same-neighborhood). Coefficients use the event-type signing in the main text; the omitted $t-1$ period is shown as an open dot at zero. Whiskers are 95% confidence intervals.

D.4 Positive-use tag breadth

Figure 8 reports signed pooled pre/post DiD coefficients for two alternative tagging choices. Panel A repeats the nuisance no-redevelopment sample across the three main control scopes. Panel B uses the narrower text-based *service-care* tag. The estimates retain the predicted negative sign for nuisance tags and positive sign for text-based *service-care* tags across all three control scopes.

Figure 8: Signed pooled DiD estimates across control scopes

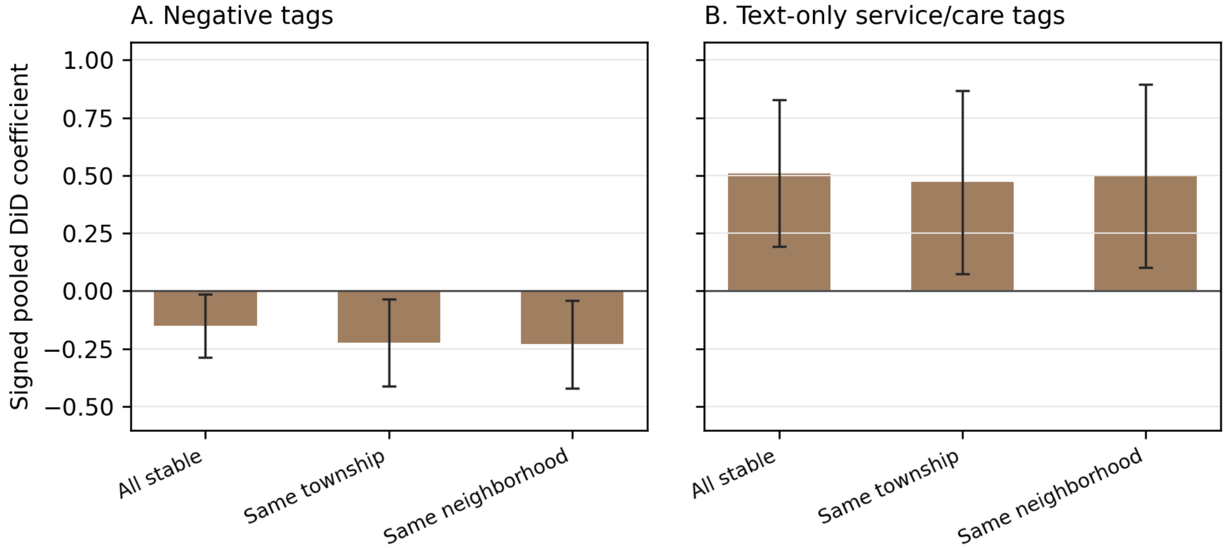

Notes: Bars report the pooled DiD coefficient comparing post years $t, t+1, t+2$ to pre years $t-2, t-1$. Coefficients use the event-type signing in the main text, so theory-consistent nuisance effects are negative and theory-consistent positive-use effects are positive. Panel A repeats the nuisance no-redevelopment sample. Panel B uses the text-based *service-care* positive-use tag, with $N = 1{,}430$ event-year observations (all-stable and same-township) and $1{,}310$ (same-neighborhood). Whiskers are 95% confidence intervals.

Figure 9 expands the positive-use definition beyond *service-care*. The enriched-core tag includes reviewed high-confidence positive-use records from permits, facilities, and business-license evidence. The enriched-broad tag additionally admits broader positive-use evidence. The top row preserves the predicted raw split-versus-merger sign pattern: splits are positive and mergers are negative. The bottom row pools splits and mergers using the event-type signing convention, and the estimates remain positive across control scopes. The broader tag increases support, but the more inclusive definition also admits noisier cases, so it serves as a breadth check rather than the main specification.

Figure 9: Positive-use tag breadth: raw and signed pooled effects

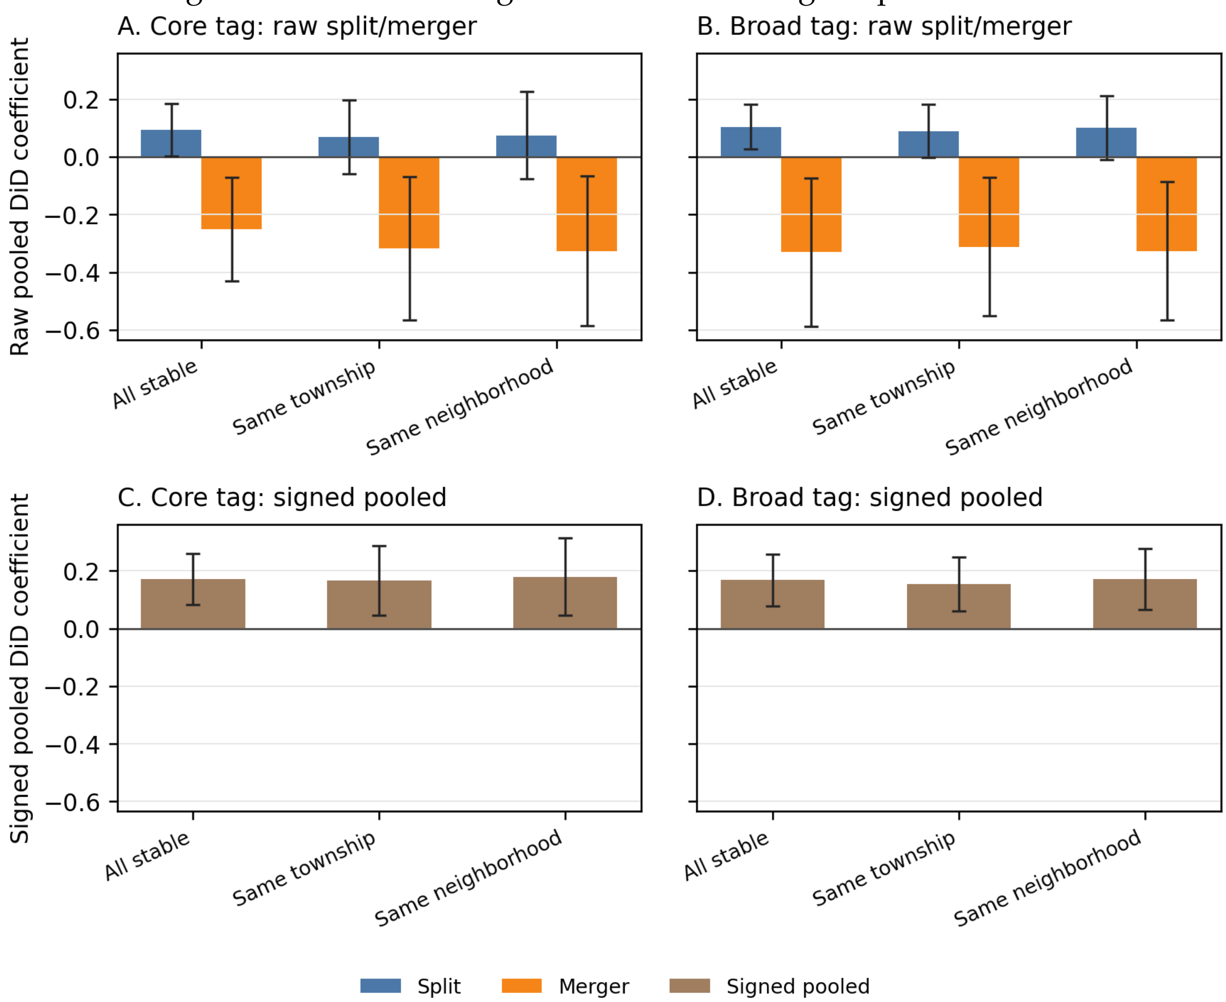

Notes: Bars report pooled DiD coefficients comparing post years $t, t+1, t+2$ to pre years $t-2, t-1$. Panels A and B show raw split and merger estimates; positive split coefficients and negative merger coefficients are theory-consistent for positive-use events. Panels C and D sign the pooled estimates by event type, leaving splits in raw scale and multiplying mergers by -1. Event-year observations: enriched-core $7{,}150$ (all-stable and same-township) and $5{,}595$ (same-neighborhood); enriched-broad $12{,}480$ (all-stable and same-township) and $9{,}890$ (same-neighborhood). The neighborhood-class-area scope is omitted to keep support comparable across panels. Whiskers are 95% confidence intervals.

References

Arnott, R. J. and Stiglitz, J. E. (1979). Aggregate land rents, expenditure on public goods, and optimal city size. *Quarterly Journal of Economics*, 93(4):471–500.

Banzhaf, H. S. and Lavery, N. (2010). Can the land tax help curb urban sprawl? evidence from growth patterns in pennsylvania. *Journal of Urban Economics*, 67(2):169–179.

Bentick, B. L. (1979). The impact of taxation and valuation practices on the timing and efficiency of land use. *Journal of Political Economy*, 87(4):859–868.

Birchwood Law (2025). How do i legally split my property into multiple lots in chicago? `https://birchwood.law/how-to-split-land-into-lots/`. Accessed May 2, 2026.

Brooks, L. and Lutz, B. (2016). From today's city to tomorrow's city: An empirical investigation of urban land assembly. *American Economic Journal: Economic Policy*, 8(3):69–105.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cellini, S. R., Ferreira, F., and Rothstein, J. (2010). The value of school facility investments: Evidence from a dynamic regression discontinuity design. *Quarterly Journal of Economics*, 125(1):215–261.

City of Chicago Department of Transportation (2024). Subdivision / dedication program application. `https://www.chicago.gov/content/dam/city/depts/cdot/permit/MapsAndPlats/2025/Application_SubdivDedicationProgram_2024.pdf`. Accessed May 2, 2026.

Clapp, J. M. and Lindenthal, T. (2022). Urban land valuation with bundled good and land residual assumptions. *Journal of Housing Economics*, 58:101872.

Cook County Assessor's Office (2026a). About the cook county assessor's office. `https://www.cookcountyassessoril.gov/about-cook-county-assessors-office`. Accessed May 2, 2026.

Cook County Assessor's Office (2026b). Divisions and consolidations. `https://www.cookcountyassessoril.gov/divisions-consolidations`. Accessed May 2, 2026.

Cook County Assessor's Office (2026c). Home value report. `https://www.cookcountyassessoril.gov/home-value-report`. Accessed May 2, 2026.

Cook County Assessor's Office (2026d). How properties are valued. `https://www.cookcountyassessoril.gov/how-properties-are-valued`. Accessed May 2, 2026.

Cook County Assessor's Office (2026e). How residential property is valued. `https://www.cookcountyassessoril.gov/how-residential-property-valued-0`. Accessed May 2, 2026.

Cook County Assessor's Office Data Department (2026). Assessor – assessed values. `https://datacatalog.cookcountyil.gov/Property-Taxation/Assessor-Assessed-Values/uzyt-m557`. Accessed May 2, 2026.

Currie, J., Davis, L., Greenstone, M., and Walker, R. (2015). Environmental health risks and housing values: Evidence from 1,600 toxic plant openings and closings. *American Economic Review*, 105(2):678–709.

de Chaisemartin, C. and D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

George, H. (1879). *Progress and Poverty: An Inquiry into the Cause of Industrial Depressions and of Increase of Want with Increase of Wealth: The Remedy*. D. Appleton and Company, New York.

Giertz, S. H., Ramezani, R., and Beron, K. J. (2021). Property tax capitalization, a case study of Dallas county. *Regional Science and Urban Economics*, 89:103680.

Haninger, K., Ma, L., and Timmins, C. (2017). The value of brownfield remediation. *Journal of the Association of Environmental and Resource Economists*, 4(1):197–241.

Hoseini, M. and Briand, O. (2020). Production efficiency and self-enforcement in value-added tax: Evidence from state-level reform in india. *Journal of Development Economics*, 144(C):None.

Hughes, C., McCluskey, W., Sayce, S., Shepherd, E., and Wyatt, P. (2018). Investigation of potential land value tax policy options for scotland: Final report. Technical report, University of Reading, for the Scottish Land Commission. Commissioned by the Scottish Land Commission.

International Association of Assessing Officers (2010). *Standard on Property Tax Policy*. Kansas City, MO.

International Association of Assessing Officers (2017). *Standard on Mass Appraisal of Real Property*. Kansas City, MO.

Linden, L. and Rockoff, J. E. (2008). Estimates of the impact of crime risk on property values from megan's law. *American Economic Review*, 98(3):1103–1127.

Lyu, X. (2024). Revisiting property tax capitalization. *Regional Science and Urban Economics*, 108:104039.

Mills, D. E. (1981). The non-neutrality of land value taxation. *National Tax Journal*, 34(1):125–129.

Mirrlees, J., Adam, S., Besley, T., Blundell, R., Bond, S., Chote, R., Gammie, M., Johnson, P., Myles, G., and Poterba, J. (2011). *Tax by Design: The Mirrlees Review*. Oxford University Press, Oxford.

Muehlenbachs, L., Spiller, E., and Timmins, C. (2015). The housing market impacts of shale gas development. *American Economic Review*, 105(12):3633–3659.

Oates, W. E. (1969). The effects of property taxes and local public spending on property values: An empirical study of tax capitalization and the Tiebout hypothesis. *Journal of Political Economy*, 77(6):957–971.

Oates, W. E. and Schwab, R. M. (2009). The simple analytics of land value taxation. In

Dye, R. F. and England, R. W., editors, *Land Value Taxation: Theory, Evidence, and Practice*, pages 51–72. Lincoln Institute of Land Policy, Cambridge, MA.

Pope, D. G. and Pope, J. C. (2015). When walmart comes to town: Always low housing prices? always? *Journal of Urban Economics*, 87:1–13.

Stiglitz, J. E. (1977). The theory of local public goods. In Feldstein, M. S. and Inman, R. P., editors, *The Economics of Public Services*, pages 84–112. Macmillan Press, London.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Tideman, T. N. (1982). A tax on land value is neutral. *National Tax Journal*, 35(1):109–111.

Titman, S. (1985). Urban land prices under uncertainty. *American Economic Review*, 75(3):505–514.

Vickrey, W. S. (2001). Site value taxes and the optimal pricing of public services. *American Journal of Economics and Sociology*, 60(5):85–96.

Wildasin, D. E. (1982). More on the neutrality of land taxation. *National Tax Journal*, 35(1):105–108.

Xing, J., Bilicka, K., Hou, X., and Raei, S. (2024). How distortive are turnover taxes? evidence from china. *Journal of Development Economics*, 171(C):None.