When Do Local Income Taxes Cause Residential Sorting?

Evidence from Pennsylvania's Piecewise Earned Income Tax

Matthew Hockert*

May 21, 2026

Abstract

Do local income taxes cause residential sorting? Pennsylvania's Earned Income Tax follows a piecewise max-rule, where a rate increase changes a worker's tax bill only when the combined home-side rate rises above the work-rate. Using block-level LEHD/LODES data from 2002–2019 and a continuous stacked difference-in-differences design, I split otherwise-identical rate increases into a binding cell that moves the worker's tax bill and a within-design placebo cell that does not. A typical binding increase reduces the count of workers who both live and work in their own municipality by roughly 13%, with clean pre-trends, while the placebo cell finding is null. I also find that rate-setters have contrasting effects on outcomes, where municipal rate increases cause negative sorting and school-district rate increases find positive sorting into a community. These findings indicating that workers sort on the bundle of local taxes and amenities rather than on statutory rates alone.

*University of Minnesota. Email: hocke069@umn.edu

1 Introduction

The Tiebout hypothesis predicts that residents sort across jurisdictions which offer different fiscal amenities at different costs and residents will 'vote with their feet' and choose specific bundles of amenities for a given amount of taxation, this is called the Tiebout margin (Tiebout, 1956). A survey of the literature conducted by Agrawal (2026) highlights the current findings of state and local income taxation and identifies sub-state (county, municipality, school districts) income-tax sorting as one of the field's open questions. The literature has focused mostly on state tax adjustments on high-earners (Agrawal & Foremny, 2019; Young et al., 2016). The research shows that top-income households are generally mobile when rates between states are large, and Agrawal (2015) documents sharp discontinuities among commodity-taxes at sub-state borders. The closest existing research comes from Yang and Heim (2017) which finds small migration effects for county-level income tax adjustments in Indiana.

This paper estimates municipal-level migration effects among workers in Pennsylvania from 2002–2019, by merging local income tax data with the LEHD Origin-Destination Employment Statistics (LODES), which links each worker's residence to their workplace at annually and at sub-municipal geography. In Pennsylvania there are roughly 2,500 municipalities and 500 school districts (SD) that set their own Earned Income Tax (EIT) rates (Agrawal, 2026; Bagchi, 2022), producing a deep cross-section of multiple sub-state income-tax opportunities. The EIT offers two additional dimensions of variation that my empirical design exploits. First, the EIT is set independently by two overlapping rate-setting jurisdictions (municipalities and school districts) and a resident's home-side tax bill is the sum of the municipal resident rate and the school-district rate. This paper's primary focus is the municipal resident rate while the school district rate enters the empirical design but is treated as secondary. Second, the EIT operates using a piecewise maximization rule where each worker's bill is $T = \max(\tau^{\text{home}}, \tau^{\text{work,muni}})$, where $\tau^{\text{home}} = \tau^{\text{res,muni}} + \tau^{\text{SD}}$ is the sum of the municipal resident rate and the school-district resident rate at the worker's address (the home stack), and $\tau^{\text{work,muni}}$ is the work rate of the worker's workplace municipality. The workplace rate is also referred to as the nonresident rate, however, I refer to it as the work rate because the primary outcome variable of is the number of people that live and work within a municipality, which is the only variable that has the cleanest causal identification to the work and home-stack rate changes A statutory home-side rate change only moves the bill if the home stack is on the binding side of the max. Act 32 of 2008 (effective 2012) consolidated EIT collection into county-based Tax Collection Districts and mandated employer withholding at the higher of the two rates, which made the max-rule mechanically enforced. Bagchi (2022) estimates that Act 32 resulted in a 14% increase in compliance and much higher local revenues. My paper does not analyze compliance but rather estimates residential-sorting and provides additional analysis pre-post Act 32.

Beyond contributing to the general literature on local tax rate changes and residential sorting this paper's methodological contribution is utilizing the piecewise structure of the EIT to generate a binding decomposition to test how rate changes that move the tax bill from the work side to the home side effect residential sorting. A statutory rate change moves the worker's bill only when the home stack is on the binding side of the maximization function, but that same statutory change produces a predicted-null when the home stack stays below the work rate. Splitting rate-change events by whether or not they are binding isolates whose bill actually moves and those whose bill does not. This within-design heterogeneity is what Agrawal (2026) identifies as missing from the existing U.S. local-income-tax literature. The decomposition is implemented at the municipality-year level in a continuous stacked difference-in-differences design that uses never-changed municipalities as the control pool. The school-district rate enters the design as a co-treatment on the home-side stack but is not used as a separate unit of analysis. Section 3 formalizes the binding decomposition and lays out the worker-bill differencing function each cell implies.

The design shows that when the municipality-level tax is increased on an already binding home-stack (Inc-bind) residents who live and work there tend to decline. At the average

Inc-bind dose of 0.24 pp cumulative over five years, the muni-rate change reduces the log count of workers who both live and work in the municipality by roughly 13% ($\beta = -0.573$ per pp, $0.24 \times -0.573 = -0.14$ log points). The increase becomes (Inc-becomes) cell, where an increase pushes a previously non-binding home stack above the work rate produces a smaller per-pp estimate of -0.250, implying cumulative dose of $0.41 \times -0.250 = -0.10$ log points (about -10%). The Inc-nobind cell is a null within-design placebo, indicating that the mechanism is producing estimates that match the hypothesized effects.

The remainder of the paper is organized as follows. Section 2 describes the institutional setting and Act 32. Section 3 introduces the six-cell binding decomposition. Section 4 describes the downloading and merging of the multiple data sources: LODES, the DCED rate panel, and the municipal and school-district controls. Section 5 explains the continuous stacked difference-in-differences design at the municipality level. Section 6 reports the pooled benchmark and the municipality-level causal headline. Section 7 reports robustness diagnostics, and Section 8 concludes.

2 Institutional Background

Pennsylvania has a deep cross-section of jurisdictions setting their own income tax rates with approximately 2,500 municipalities and 500 school districts each levy a separate Earned Income Tax on essentially all earned income earned in or by Pennsylvania residents (Agrawal, 2026; Bagchi, 2022). Unlike property taxes, which are paid by owners and capitalized into home values, the EIT falls directly on workers and varies across municipal boundaries, creating sharp spatial discontinuities in tax burden that can influence where workers choose to live and work.

2.1 Pennsylvania Earned Income Tax

Pennsylvania's EIT is mandatory for all municipalities under the Local Tax Enabling Act of 1965 (Act 511; Pennsylvania General Assembly, 1965) and applies to all earned income. The standard combined cap is 1%, typically split as 0.5% municipal and 0.5% school district. Home Rule municipalities may exceed this cap (Pennsylvania General Assembly, 1972), and Philadelphia operates separately under the Sterling Act of 1932 (Pennsylvania General Assembly, 1932) at approximately 3.9%. EIT is a valuable revenue funder for local governments next to property-taxes. However, quantifying the effect of property tax millage on sorting would require a 2002–2019 municipality-level millage panel that is not currently available from DCED, and I leave the comparison to property-tax sorting as a separate exercise. Out of caution that property taxes may impact sorting I control for the share of revenue that a municipality collects from property taxes, EIT, and transfers within a given year.

Each Pennsylvania municipality has three EIT rates relevant to a worker: the municipality resident rate $\tau^{\text{res,muni}}$ levied on residents of the municipality, the municipality work rate $\tau^{\text{work,muni}}$ levied on workers employed in the municipality, and the school-district rate τ^{SD} levied on residents only (school districts cannot tax nonresidents). Each address is associated with a six-digit Political Subdivision (PSD) code identifying its (municipality, school district) pair. Workers report this code on the Residency Certification Form, and employers use it to determine withholding obligations.

Under Act 32 (effective January 1, 2012), a worker's total EIT bill is determined by the maximum of the home-side rate stack and the work-location rate, and the bill is split piecewise between home and work jurisdictions:

$$\left(\tau^{\text{home}},\ \tau^{\text{work}}\right) = \begin{cases} \left(\tau_j^{\text{res,muni}} + \tau_j^{\text{SD}} - \tau_i^{\text{work,muni}},\ \tau_i^{\text{work,muni}}\right) & \text{if } \tau_j^{\text{res,muni}} + \tau_j^{\text{SD}} > \tau_i^{\text{work,muni}} \\ \left(0,\ \tau_i^{\text{work,muni}}\right) & \text{if } \tau_j^{\text{res,muni}} + \tau_j^{\text{SD}} \leq \tau_i^{\text{work,muni}} \end{cases} \tag{1}$$

where j is the worker's home municipality and i is the work municipality. The work jurisdiction always collects its full work rate first, then the home municipality and school district split whatever portion of the worker's bill is left over after the work claim. If the work rate exceeds the home stack, the home jurisdictions get nothing. The worker's total bill is the higher of the two stacks:

$$\mathrm{Bill}_{ji} = \max\bigl(\tau_j^{\mathrm{res,muni}} + \tau_j^{\mathrm{SD}},\ \tau_i^{\mathrm{work,muni}}\bigr). \tag{2}$$

Equation (1) implies which margin a worker can use to respond to a rate increase. When the home stack binds ($\tau_j^{\mathrm{home}} > \tau_i^{\mathrm{work,muni}}$), the worker can lower the bill only by moving their residence to a lower-rate municipality or school district. When the work rate binds, the worker can lower the bill only by changing jobs to a lower-rate work location, while moving homes without changing jobs has no effect. For workers who live and work in the same municipality, the home stack and work rate are both set by that municipality, and the bill is determined by the home stack on the binding side.

A worker who lives in municipality j and moves to a lower-rate municipality j' while keeping the same job at i saves at least the SD rate at j: the new bill is $\max(\tau_{j'}^{\mathrm{home}}, \tau_i^{\mathrm{work,muni}})$, which is at most the old bill τ_j^{home} minus τ_j^{SD}. The empirical design follows this logic directly by using both the treatment $\Delta\tau^{\mathrm{muni}}$ and the binding classification built from the worker's home stack. The regressions therefore measure sorting in response to the same tax a mover would actually face. This paper does not estimate the effect of nonresident (workplace tax) increases on sorting due to the lack of binding rate changes of the workplace tax over the home-stack. For the work side to bind, a single rate-setter must exceed the home-stack, which under the standard Act 511 rate structure almost never occurs: across 2002–2019, the work rate is the binding side in roughly 0.2% of municipality-years. The handful of exceptions are municipalities with statutory authority to exceed the Act 511 cap (Home Rule and Act 47 distressed cities). Because there are few work-binding events, a work-side binding decomposition is not identifiable.

Philadelphia is the one exception to the Act 32 framework. It operates under the Sterling Act of 1932 where Philadelphia residents pay the city resident rate regardless of work location with no offset and Philadelphia nonresidents pay the city work rate, with no residual flowing to home jurisdictions in other Pennsylvania municipalities. My main estimates retain Philadelphia in the sample for completeness and the results are robust to dropping it.

Act 32 transformed the system by consolidating collection into approximately 67 county-based Tax Collection Districts (TCDs), each with a single appointed collector. The law mandated standardized employer withholding at the higher of the employee's resident and work rate. This had two consequences. First, it made the EIT mechanically binding where the rate changes now appear in every worker's paycheck through employer withholding, regardless of whether the worker resides in the taxing municipality. Second, it eliminated the possibility of avoidance of the work-location claim through weak enforcement. After Act 32, a worker could no longer effectively avoid the work-rate portion owed to the work municipality.

However, Act 32 did not eliminate all tax savings from residential relocation. The piecewise rule in Equation (1) continues to imply that a worker whose original home stack binds can reduce the bill by moving to a sufficiently lower-stack home jurisdiction where the new bill is $\max(\tau_{j'}^{\text{home}}, \tau_i^{\text{work,muni}})$, which falls below the old home stack τ_j^{home} whenever $\tau_{j'}^{\text{home}} < \tau_j^{\text{home}}$ and the original home stack was the binding side. Act 32 thus changed enforcement and made the max-rule mechanically binding through withholding; it did not eliminate the residential-relocation margin that the mechanism in this paper exploits. Bagchi (2022) highlighted the impact of Act 32 by showing that consolidated employer withholding increased PA EIT compliance thus raising reported earned income in the rate-raising municipality. The present paper complements that work by isolating the residential-sorting response to the same withholding-enforcement mechanism exploiting the variation in whether a given rate change actually moves the worker's bill (the binding decomposition); the pre/post-Act-32 split in Section 7 shows the binding-cell sorting response is strengthened (the effect existed prior to Act 32) after Act 32, complementing the compliance-response finding.

Figure 1 documents the institutional break at Act 32. Before 2012, fewer than 60% of PA municipalities reported a work rate to DCED; after 2012, when consolidated withholding made the max-rule mechanically binding, coverage rises sharply. This is the pre-/post-Act-32 enforcement break used in the heterogeneity test of Section 7.1.

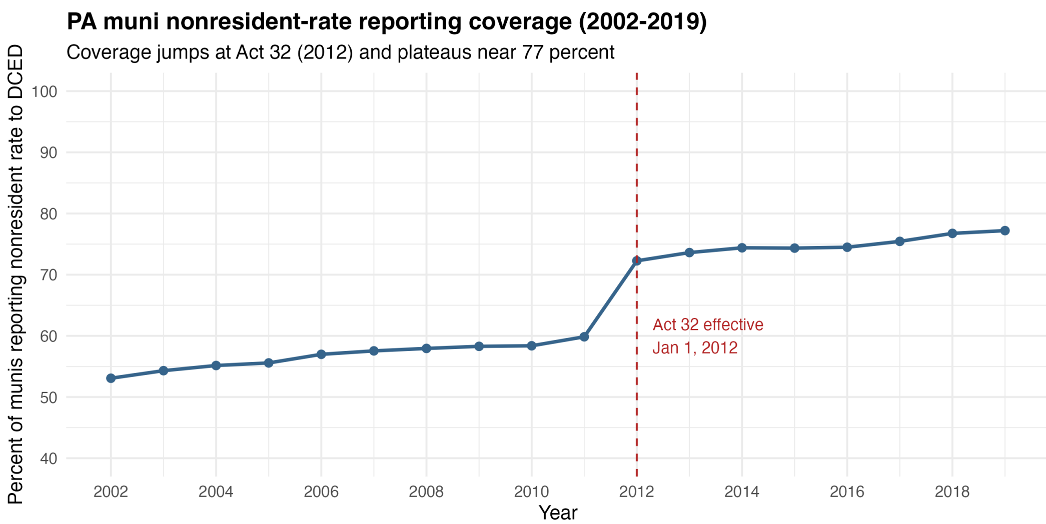

Figure 1: Share of PA municipalities reporting a work rate to DCED, by year. The break at 2012 reflects the Act 32 requirement that employers withhold at the higher of the home resident and work rates.

School districts in Pennsylvania also levy an EIT under Act 511 (Pennsylvania General Assembly, 1965), typically at 0.5% (the other half of the standard 1% cap). Act 1 of 2006 (Pennsylvania General Assembly, 2006) (the Taxpayer Relief Act) additionally authorized school districts to exceed the 0.5% cap specifically to fund property tax relief. I find that the correlation between the municipal and school-district resident rates is only weakly negative of -0.12 indicating the two taxing jurisdictions move on relatively separate schedules. Since every census block sits in both a municipality and a school district, workers face two overlapping EIT rates that can change independently, providing a natural decomposition of the residential sorting response into municipal and school district components.

3 Conceptual Framework: Binding Decomposition

The piecewise tax allocation in Equation (1) is the causal identification backbone of my empirical design. A worker's bill is $T = \max(\tau^{\text{home}}, \tau^{\text{work,muni}})$. A rate change only moves the bill when its side of the max binds. The implied change in the worker's bill from a single rate event is

$$\Delta T \;=\; \max(\tau^{\text{home}}_{\text{post}}, \tau^{\text{work,muni}}) - \max(\tau^{\text{home}}_{\text{pre}}, \tau^{\text{work,muni}}),$$

which depends jointly on the direction of the rate change and on which side of the max is binding pre- and post-event. This tax difference from pre and post rate changes can produce a strong sorting response, a moderate response, or no response at all depending on the binding transition.

3.1 Six Cells per Rate-Setter

Two definitions structure the decomposition. First, the piecewise rule $T = \max(\tau^{\text{home}}, \tau^{\text{work,muni}})$ has two sides: the home side $\tau^{\text{home}} = \tau^{\text{res,muni}} + \tau^{\text{SD}}$, set at the worker's residence, and the work side $\tau^{\text{work,muni}}$, the work rate set by the worker's workplace municipality. Whichever side is larger is the binding side and determines the bill. Second, three rate-setters adjust statutory rates independently across municipality-years: the municipality resident rate $\tau^{\text{res,muni}}$, the school-district rate τ^{SD}, and the municipality work rate $\tau^{\text{work,muni}}$. The decomposition is built one rate-setter at a time.

For each rate-setter τ, a cell classifies a municipality by the direction of its first rate event in τ (Increase or Decrease) and then the binding transition of its home stack across that event: stayed-binding, flipped, or stayed-non-binding. The cross gives six cells plus a seventh never-changed control pool, of which only the three Increase cells and the NC control are used in this paper (Table 1). Each municipality therefore carries three cell labels (one for each of $\tau^{\text{res,muni}}$, τ^{SD}, and $\tau^{\text{work,muni}}$) and the binding state for each is determined by where its home stack sits relative to its work rate at the time of that event.

Cell	Direction	Pre-binding?	Post-binding?	Bill effect
Inc-bind	$\uparrow$	yes	yes	rises by full $\Delta\tau$
Inc-becomes	$\uparrow$	no	yes	rises by part of $\Delta\tau$
Inc-nobind	$\uparrow$	no	no	no change (placebo)
NC	none	—	—	control pool

Table 1: Binding cells per rate-setter on the increase side. For Inc-becomes the bill rises only by the portion of $\Delta\tau$ above the work rate. Three mirror decrease cells exist by construction but are not estimated in this paper because the decrease cohort is too thin to support causal inference (5 Dec-bind, 3 Dec-becomes-non, 11 Dec-nobind) and the realized-dose pre-trend rejects parallel trends on every Dec outcome.

The cell labels describe the state of the home stack at the rate event and not necessarily which rate-setter is caused the cell to be that state. The home stack $\tau^{\text{home}} = \tau^{\text{res,muni}} + \tau^{\text{SD}}$ is additive in its two home-side components where either component can put the stack above or below the work floor. A municipality therefore carries two independent cell labels, one for each rate-setter (municipality and school district), and the two are not mutually exclusive such as when a muni classified as Inc-bind on the muni side may be Inc-becomes on the SD side (or any other combination). The interaction terms in the estimating equation (Section 5) separate the per-pp effect of a muni-rate change from the per-pp effect of an SD-rate change, each conditional on its own cell label.

What the labels imply for an individual worker depends on where that worker is employed. Under Act 32 the employer withholds at $\max(\tau^{\text{home}}, \tau^{\text{work,muni}})$, where $\tau^{\text{work,muni}}$ is the work rate of the worker's workplace (Bagchi, 2022; Pennsylvania General Assembly, 2008). For a worker who lives and works in municipality i the workplace rate is exactly i's own $\tau_i^{\text{work,muni}}$, which is the comparison floor the cell classification uses. For a resident of i who works in a different municipality j, the relevant floor is $\tau_j^{\text{work,muni}}$, which the LODES origin–destination data does not let us recover at the worker level and therefore the label using i's own work rate is an aggregate proxy that is reasonable on average across Pennsylvania's local commute shed but not verifiable individually. This distinction motivates the outcome ordering in Section 6:

the count of workers who live and work in their own municipality, log(live-and-work), is the cleanest binding-rule test because the classification matches each worker's actual workplace and the model directly estimates the response of those who face either τ^{home} or $\tau^{\text{work,muni}}$. The other count outcomes are noisier tests: log(total jobs) and the share of workers living in their own municipality both fold in in-commuters, and log(in-commuters) does not test the muni resident-rate binding rule at all because in-commuters' bills do not depend on the resident rate of the municipality where they work.

The fundamental testable prediction is an ordered ranking of cell magnitudes where Inc-bind effects exceed Inc-becomes effects, which in turn exceed Inc-nobind effects, with the last close to zero and serving as the within-design placebo. A tax salience argument would state that workers in the Inc-becomes cell experience a discrete switch from paying zero home-side tax to paying a positive amount, which should trigger a large behavioral response. The central empirical comparison of the paper is between Inc-bind and Inc-nobind where if statutory rates matter mechanically through tax liability, the residential response should appear in Inc-bind and not in Inc-nobind. If Inc-nobind moves, the mechanism interpretation fails and the response is something other than the binding rule because Inc-nobind events do not change the worker's bill. This makes the design self-falsifying in a way that a conventional pooled-Increase specification is not. In the pooled specification, every increase is a treated event, while here the institutional rule predicts that only a subset of statutory increases should generate a response.

4 Data

4.1 Data Sources

The analysis merges the LEHD residence–workplace dataset and the Pennsylvania EIT rate records to create a panel of rate changes and workforce data over space and time. The two datasets do not share common identifiers, and was bridged using TIGER 2020 block geometry, which nests blocks within county subdivision code (cousubs). A municipality-name crosswalk maps each DCED municipality ID to its corresponding COUSUBFP code, matching 2,496 of 2,505 PA municipalities (99.6%). The remaining 9 are small boroughs that have either dissolved or merged during 2002–2019 and are dropped from the analysis.

The residence–workplace data come from LEHD LODES Version 8, 2002–2019 (U.S. Census Bureau, 2024b). The Census Bureau applies synthetic noise infusion at the block level to protect respondent confidentiality and is calibrated so that aggregation to tract and larger geographies approximates ground truth (Graham et al., 2014). The sample period of 2002–2019 reflects LODES coverage beginning in 2002 and 2020–2023 to not include disruptions in the data resulting from the COVID Pandemic.

Municipal and school-district EIT rates come from the Pennsylvania Department of Community and Economic Development (DCED) MunStats portal (Pennsylvania Department of Community and Economic Development, 2024a), scraped annually. In some DCED reports, the flat-dollar Local Services Tax is filled in the same field as the EIT in some years, so rates above 5% are treated as data-entry errors and dropped except for Home Rule cities and Philadelphia.

Municipal fiscal controls come from the PA DCED Statewide Municipal Annual Financial Report database (Pennsylvania Department of Community and Economic Development, 2024b). School-district fiscal controls come from the Census Annual Survey of School System Finances (F-33) (U.S. Census Bureau, 2024a); the NCES `LEAID`-to-DCED SD ID crosswalk covers all 500 school districts. All fiscal controls are share-based and time-varying. Municipality regressions use six DCED AFR shares: three expenditure shares (police-services, streets, debt-service) and three revenue shares (EIT-revenue, real-estate-tax, intergovernmental-revenue). LODES blocks are aggregated to muni for analysis.

The municipality–SD assignment uses the modal SD ID per muni in the DCED rate panel (97.6% of municipalities are single-SD), so each muni-year carries a single τ^{SD} co-treatment. Municipality regressions further impose a minimum employment threshold of 50 total jobs, which protects the share-based outcomes from mechanical noise (a single worker moving changes a share by $1/N$, so at $N = 10$ a one-worker move dominates the per-pp coefficient) and matches the LODES synthetic-noise calibration that is validated at tract scale and above (Graham et al., 2014). The threshold drops 185 of 2,559 municipalities (7.2%), almost all sparsely populated NC-pool townships; the treated cohort is essentially unaffected.

4.2 Descriptive Statistics

This section presents the summary statistics at the municipality level. The primary outcome variables are resident-worker counts, workplace counts, in-commuter flows, and the live-and-work share. Fiscal outcomes used as control variables (log EIT revenue, log total tax revenue, log total revenue, log police expenditure) are reported in Appendix C.

Table 2 summarizes the estimation panel.

Table 2: Muni-year summary statistics (2002–2019)

Variable	N	Mean	SD	Min	Max
Resident rate τ^M (%)	43,374	0.525	0.163	0.000	4.331
Nonresident rate $\tau^{nr,M}$ (%)	34,080	0.883	0.237	0.000	3.820
SD rate τ^{SD} (%)	43,374	0.622	0.289	0.000	2.050
Home stack $\tau^M + \tau^{SD}$ (%)	43,374	1.147	0.316	0.000	4.331
log live-and-work	40,748	3.638	2.011	0.000	12.877
log total jobs	45,291	5.819	2.006	0.000	13.353
Share live-and-work in muni	45,291	0.095	0.081	0.000	1.000

Notes: Muni-year level; one observation per (municipality, year). Rates in percentage points. $\tau^{nr,M}$ has reporting gaps before Act 32; the N for that row is smaller than the full panel.

Table 2 presents the municipality level summary statistics. Table 2 shows that there are roughly 43,374 muni-year observations, and the resident rate is tightly clustered near the Act 511 default at $\bar{\tau}^M = 0.53\%$ and a standard deviation of 0.16pp. There is a long right tail indicated by the max rate of 4.33% concentrated in Philadelphia and a handful of Home-Rule and Act 47 distressed cities. The work rate has 34,080 observation, about 9,300 fewer observations than the resident rate, due to DCED's lack of coverage for the work-rate field being incomplete before Act 32 mandated reporting. The SD rate has nearly twice the muni-rate dispersion (SD 0.29pp vs. 0.16pp) and maxes at 2.05%, potentially making the school-district rate-setter a larger statistical lever. The home stack averages 1.15%, above the 1% Act 511 cap, because Act 1 SD increases and Home-Rule muni rates push the upper level past the ceiling. Only 9.5% of muni jobs are held by residents of the same muni showing that residential sorting operates on a thin within-muni margin.

Table 3: Cell counts: muni resident-rate side, 2002–2019

Cell	Munis	Muni-years	Mean dose (pp)	Extreme dose (pp)
Never-changers	1,687	30,358	0.000	0.000
Inc-bind	24	432	0.168	0.850
Inc-becomes	39	702	0.225	1.250
Inc-nobind	14	252	0.301	0.650

Notes: Cells classify each muni's first resident-rate event by increasing and the binding-state transition relative to the muni's nonresident rate. NC is the never-changed control pool. Dose is the cumulative $\Delta\tau$ within the 5-year event window. The Extreme dose column reports the max cumulative dose for Inc cells.

Table 3 reports the cell assignments on the muni resident-rate side. Of roughly 2,560 municipalities, after dropping those that decrease or flip-flop any of their three rates, 1,687 never-changers form the NC control pool and only 77 enter one of the three increase cells, split as 24 Inc-bind, 39 Inc-becomes, and 14 Inc-nobind. Inc-becomes shows a larger average dose than Inc-bind, indicating that crossing the binding threshold typically requires a bigger statutory move than slowly inching up an already-binding rate.

Figure 2 reports the timing of municipal resident-rate changes across the panel. Increases jumped in the early 2000s, from 2002 to 2009 and taper off after. Most of the identifying variation in my models predates Act 32, highlighting the significant effects found are on smaller increases and further warranting a pre-post 2012 decomposition.

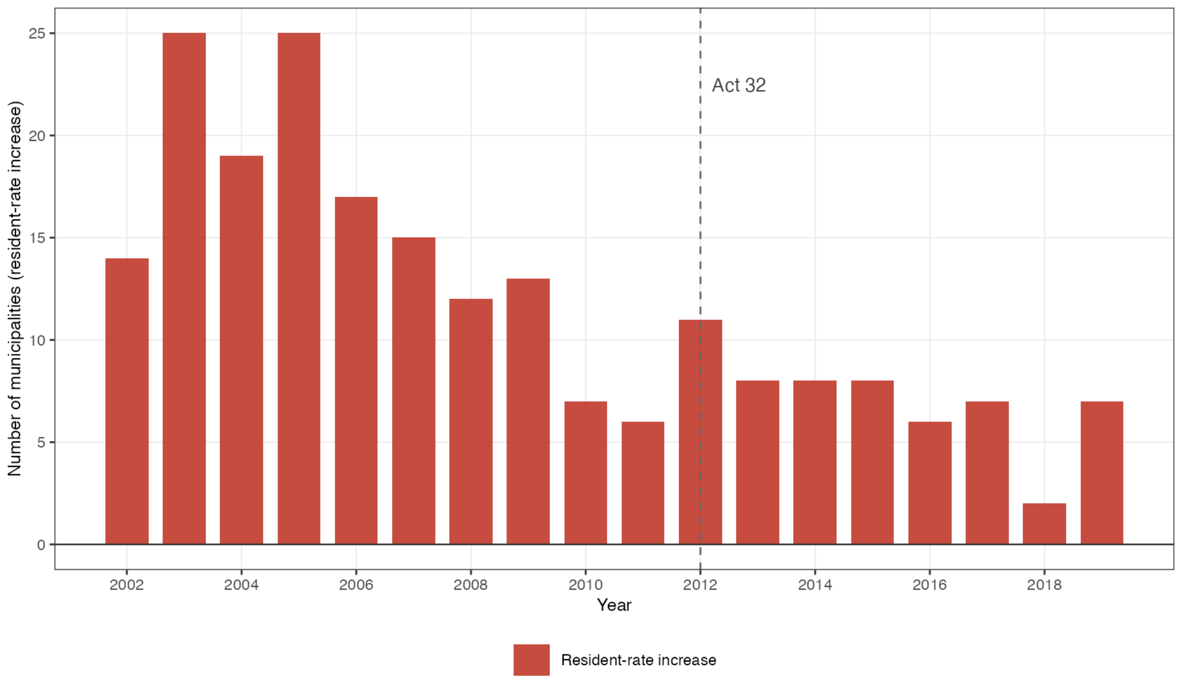

Figure 2: Number of PA municipalities changing their resident rate, by year. Increases shown above zero; decreases shown below. Dashed line: 2012, when Act 32 consolidated withholding.

Lastly, figure 3 shows the spatial distribution of the rate changers in my sample. Most rate changers occur along the east coast, but some are also scattered within the interior of the state. This spatial distribution highlights the value of future work to understand the effects when two neighboring municipalities raise or lower rates across time.

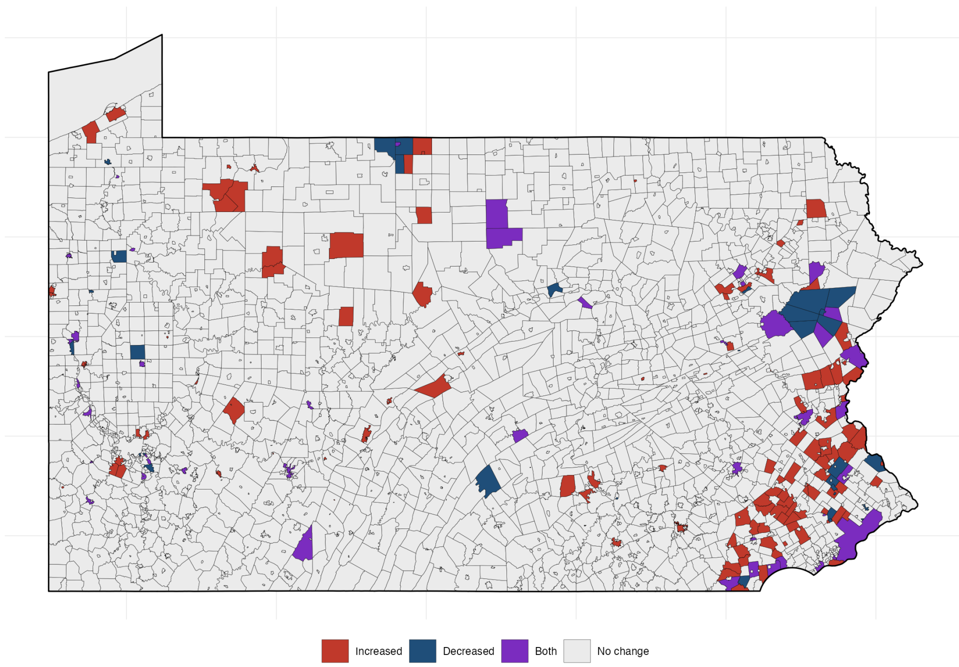

Figure 3: Pennsylvania EIT rate changes, 2000–2019. Red = increased; blue = decreased; purple = both; grey = no change.

5 Empirical Strategy

I estimate two specifications using a continuous stacked difference-in-differences design. First, I run a pooled-Increase aggregate (§5.1) that collapses all rate-increase events at each rate-setter into a single only-increase group. Second, I estimate a binding decomposition (§5.2) on the three Increase cells defined in Section 3, isolating the rate changes that move the worker's bill from those that do not. Both specifications are estimated at the municipality-year level. Decrease cells are not estimated: the cohort is thin (5 Dec-bind, 3 Dec-becomes-non, 11 Dec-nobind) and the realized-dose pre-trend rejects parallel trends on every Dec outcome (Wald $p < 0.001$), so the Dec coefficients cannot be interpreted as causal. It is important to note that a theoretical case for symmetric responses to rate cuts and rate hikes is important but most likely weak. First, moving costs are sunk, so the effect of rate cuts may be less salient to incumbent residents than rate hikes. Pre-trends are tested separately by a realized-dose Sun–Abraham event study (§5.3). This section focuses on estimating equations, identification, and diagnostics.

5.1 Pooled Aggregate Specification

In this section I report an aggregate specification that pools the treated Increase cells per rate-setter into a single directional group Inc (any rate increase), with never-changers (NC) as the omitted reference. Decrease cells are excluded because the cohort is thin (5 Dec-bind, 3 Dec-becomes-non, 11 Dec-nobind) and the realized-dose pre-trend rejects parallel trends on every Dec outcome (Wald $p < 0.001$). The pooled regression answers: what is the average per-percentage-point (pp) effect of a rate hike on residents, ignoring whether or not the hike binds? The binding decomposition that follows in Equation (4) then explains how the binding mechanism drives the pooled coefficient.

For each unit-time observation (i, t) on the stacked panel, the per-pp pooled regression is

$$\begin{aligned} Y_{\tilde{i}\tilde{t}} \;=\; \alpha_{\tilde{i}} \;+\; \gamma_{\tilde{t}} + \beta^{M}_{\text{Inc}}\, \Delta\tau^{M}_{it}\, \mathbf{1}\{d^{M}_{i} = \text{Inc}\} \\ + \beta^{SD}_{\text{Inc}}\, \Delta\tau^{SD}_{it}\, \mathbf{1}\{d^{SD}_{i} = \text{Inc}\} \\ + \mathbf{X}'_{it}\boldsymbol{\phi} \;+\; \varepsilon_{\tilde{i}\tilde{t}}. \end{aligned} \tag{3}$$

The dependent variable $Y_{\tilde{i}\tilde{t}}$ is the outcome of interest. At the municipality level the outcomes are the logs of live-and-work workers (residents employed in their own municipality), in-commuters, log total jobs, and the share of workers living in their own municipality. The unit i is a municipality and t a calendar year. The subscripts $\tilde{i} = (i, c)$ and $\tilde{t} = (t, c)$ are the unit-by-stack and year-by-stack groupings where each treated unit i appears in exactly one sub-experiment (its first-event cohort c), while each NC unit appears in every cohort it overlaps. Cohort c is defined as the first year that the rate-setter changed τ. The fixed effects $\alpha_{\tilde{i}}$ and $\gamma_{\tilde{t}}$ are for each unit and cohort pair. The directional grouping $d^{s}_{i} \in \{\text{NC}, \text{Inc}\}$ classifies unit i's rate trajectory on rate-setter $s \in \{M, SD\}$, with NC the omitted reference as they are the control group. The dose $\Delta\tau^s_{it} = \tau^s_{it} - \tau^s_{i,c-1}$ is the cumulative change in the rate from the pre-event baseline ($c-1$, one year before the first event) to year t, measured in percentage points; it is zero in pre-event years and for NC units, and accumulates with subsequent rate changes within i. The control vector $\mathbf{X}_{it}$ holds the six municipality fiscal share controls (police-services, streets, debt-service, EIT-revenue, real-estate-tax, intergovernmental-revenue). $\boldsymbol{\phi}$ is the coefficient vector on these controls and $\varepsilon_{\tilde{i}\tilde{t}}$ is the error term.

5.2 Binding Decomposition: Stacked Continuous DiD

The binding decomposition is the main estimator of the paper. It replaces the two-level direction grouping d^s_i from Equation 3 with a multilevel cell cell^s_i that entails both the direction of unit i's first event and whether the home stack was binding pre- and post-event (Section 3.1, Table 1). This splits the increase side into Inc-bind, Inc-becomes, and Inc-nobind, while NC remains the omitted reference group. Decrease cells are not estimated for the reasons discussed in 5. The stacked panel, unit-by-stack groupings, dose $\Delta\tau^s_{it}$, controls $\mathbf{X}_{it}$, fixed effects, and policy-jurisdiction clustering are identical to Equation 3. The sub-experiment for each cohort c runs on the window $t \in [c-K, c+K]$ with $K=5$, which gives the realized-dose ES diagnostic four pre-event leads (Section 5.3) and covers the bulk of the treated units with a full 5-pre, 5-post window.

For each rate-setter s and outcome Y the stacked-continuous regression is

$$\begin{aligned} Y_{\tilde{i}\tilde{t}} \;=\; \alpha_{\tilde{i}} \;+\; \gamma_{\tilde{t}} + &\sum_{k\in\mathcal{C}_{-\text{NC}}} \beta^{\text{muni}}_k \,\Delta\tau^{\text{muni}}_{it} \cdot \mathbf{1}\{\text{cell}^{\text{muni}}_i = k\} \\ &+ \sum_{k\in\mathcal{C}_{-\text{NC}}} \beta^{\text{SD}}_k \,\Delta\tau^{\text{SD}}_{it} \cdot \mathbf{1}\{\text{cell}^{\text{SD}}_i = k\} \\ &+ \mathbf{X}'_{it}\boldsymbol{\phi} \;+\; \varepsilon_{\tilde{i}\tilde{t}}. \end{aligned} \tag{4}$$

The set $\mathcal{C}_{-\text{NC}} = \{\text{Inc-bind, Inc-becomes, Inc-nobind}\}$ collects the three estimated treated cells; cell_i^s is the cell label on rate-setter s and the unit level being modeled i. The regression returns three β_k^s coefficients per rate-setter, each the per-pp dose-response of Y to a 1 pp change in rate s restricted to units in cell k. The $\beta_{\text{Inc-nobind}}^{\text{muni}}$ coefficient is the within-design placebo meaning that the rate hike does not move the worker's bill in this cell, so the coefficient should be statistically indistinguishable from zero. All other elements of the regression are inherited from Equation 3.

The construction uses the negative-weighting fix from Baker et al. (2025): treated units enter exactly one stack (their first-event cohort), NC units enter every stack, so no treated unit serves as control to a later-treated unit. The cell decomposition is a novel addition because it generates a ladder of predictable effects including the within-design placebo (Inc-nobind), which the pooled stacked DiD classification cannot produce.

The identifying variation in this design comes from the cross-unit cumulative rate changes after each unit's first event. Figure 4 and Figure 5 show the average change post event by rate-setter. Each line plots the cross-muni mean and inter-quartile range of cumulative $\Delta\tau$ at each event-time. Figure 4 traces the mean cumulative resident-rate change of Inc-bind units relative to their pre-event baseline. The pre-event window is mechanically flat at zero by construction (dose is measured relative to the unit's own $c - 1$ rate) and the post-event window shows the realized cumulative $\Delta\tau^{\text{muni}}$ rising to roughly 0.24 percentage points by $k = +5$. While the school district rate sits far above the municipal rate at roughly 0.60 percentage points.

Figure 4: Cumulative Inc-bind dose by event-time: municipality cohort vs. school-district cohort.

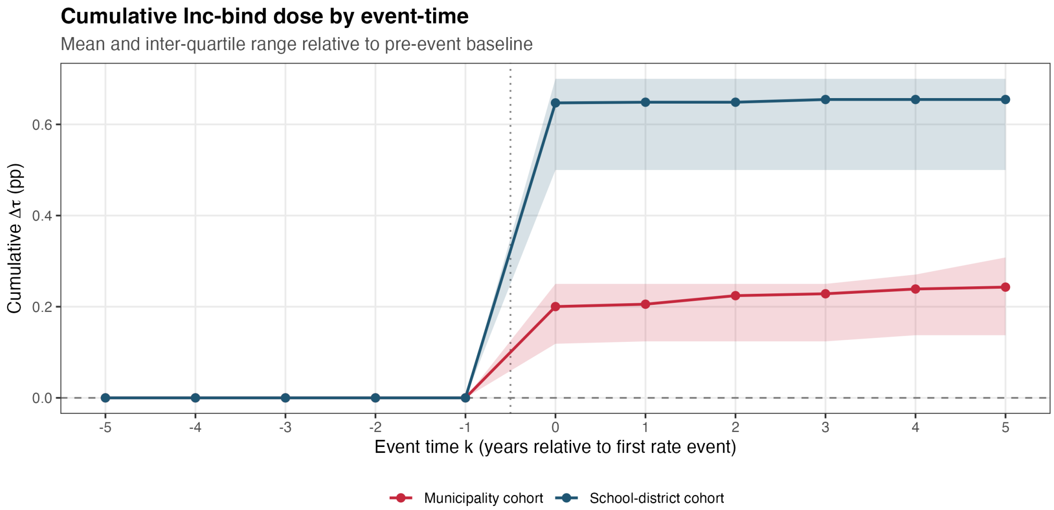

Notes: Mean and inter-quartile range of cumulative dose by relative event-time k. The municipality cohort (red) reaches a mean of 0.24 pp at $k = +5$; the school-district cohort (blue) reaches 0.65 pp. Pre-event values are zero by construction.

Figure 5 traces the mean cumulative resident-rate change of Inc-becomes binding on units relative to their pre-event baseline. The post-event window shows a similar effect as figure 4 where the school district rates sit above municipality rates at 0.60 percentage points but the municipality rate above 0.40 which is a much larger increase than the inc-binds grouping thus shrinking the gap between the average school district and average municipality rate. Figure 5 indicates that for an the rate to become binding the jump may primarily due to municipality rate adjustments.

Figure 5: Cumulative Inc-becomes dose by event-time: municipality cohort vs. school-district cohort.

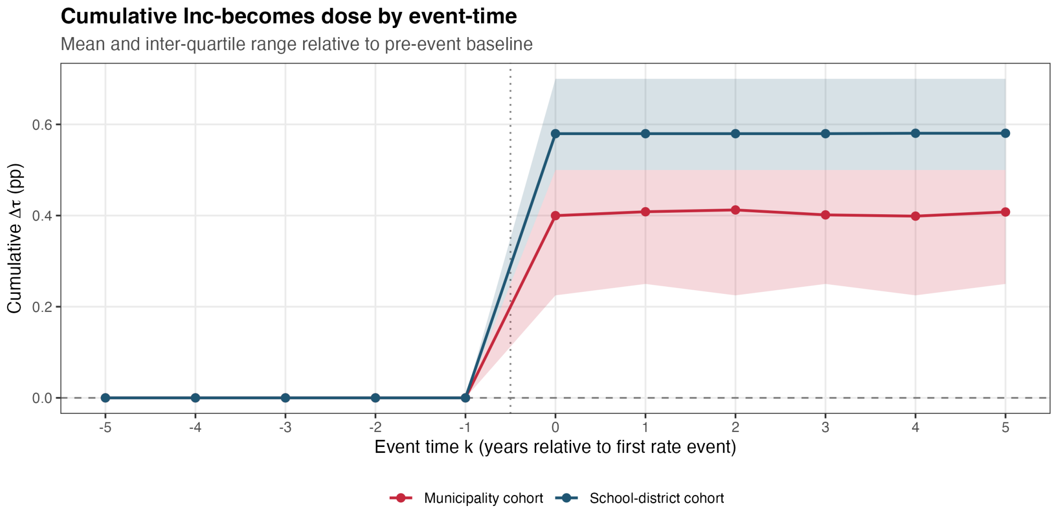

Notes: As in Figure 4 but for the Inc-becomes cell, where the home stack flips from non-binding to binding at the first event.

5.3 Realized-Dose Event-Study Diagnostic

Equations 3 and 4 cannot test for pre-treatment leads directly because $\Delta\tau^s_{it}$ is mechanically zero for all $t < c$, so any pre-event coefficient on the dose is also mechanically zero. I therefore run a separate realized-dose Sun–Abraham event study on the same stacked panel to test parallel trends, the condition recommended for continuous-treatment DiD designs by Callaway and Sant'Anna (2021) and de Chaisemartin et al. (2024) states that units with larger eventual rate changes should not be on differentially sloped pre-event trajectories relative to NC controls. For each cell on rate-setter s I restrict the panel to units in (NC, cell) and run

$$Y_{it} = \alpha_{\tilde{i}} + \gamma_{\tilde{t}} + \sum_{k \neq -1} \beta^M_k \, \mathbf{1}\{\ell_{it} = k\} \, \widetilde{D}_i + \mathbf{X}'_{it}\boldsymbol{\phi} + \varepsilon_{it}. \tag{5}$$

Two variables are new relative to Equation 4. The event-time index $\ell_{it} = t - c^M_i$ counts the years between calendar year t and unit i's first muni-rate event year c^M_i, with $\ell = -1$

(one year before the event) as the omitted reference; in the sum below k denotes a value of ℓ, i.e. an event-time, in contrast to Equation 4, where k indexed the binding cells. The realized-dose tag $\widetilde{D}_i = \tau^M_{i,c+5} - \tau^M_{i,c-1}$ is a single number per unit where i's eventual cumulative rate change is the observed rate five years after its first event minus the rate the year before and held fixed across all event-time years rather than allowed to grow from zero as $\Delta\tau^M_{it}$ does in Equation 4. It enters the regression only through the interaction $\mathbf{1}\{\ell_{it} = k\}\,\widetilde{D}_i$, which equals $\widetilde{D}_i$ in the single year unit i is at event-time k and zero in all its other years because that varies within unit, the unit fixed effect (which removes only within-unit constants) does not absorb it. Because the tag is non-zero in every event year and varies across units by eventual treatment intensity, each pre-event β^M_k tests whether units that eventually take larger rate changes were already on differently sloped trajectories. The time-varying $\Delta\tau^M_{it}$ cannot do this: it is mechanically zero before the event, so its pre-event coefficients are mechanically zero. The outcome Y_{it}, the unit-by-stack and year-by-stack fixed effects, the control vector $\mathbf{X}_{it}$, the units of observation, and the policy-jurisdiction clustering are the same as in Equation 4.

The pre-event coefficients β^M_k for $k \leq -2$ are the parallel-trends test: if units that eventually take large rate changes are on differentially sloped trajectories before the rate change, the β^M_k at pre-event leads will be non-zero. Under strong parallel trends they are zero. The post-event coefficients β^M_k for $k \geq 0$ recover the per-pp dose-response at each event-time and average to the corresponding cell coefficient β^s_{cell} from Equation 4 when the unit's rate path is a single step at the event year. I report the joint Wald p-value as a single pre-trend statistic for each cell-outcome combination.

Section 6 reports the headline binding decomposition at each estimating unit and the pooled-Increase benchmark. Additional diagnostics (pre/post Act 32 split, Act 47 exclusion) are in Section 7.

6 Results

I report results in two layers. First, Section 6.1 presents the pooled Inc benchmark, which approximates what a conventional local-tax design would recover by averaging across all rate-increase events without distinguishing whether they bind. Second, Section 6.2 reports the binding decomposition at the municipality level. The central empirical comparison is between the Inc-bind cell (where the statutory rate change moves the worker's bill) and the Inc-nobind cell (where it does not), with Inc-nobind serving as the within-design placebo. Dec cells are not estimated: the cohort is thin (5 Dec-bind, 3 Dec-becomes-non, 11 Dec-nobind) and the realized-dose pre-trend rejects parallel trends on every Dec outcome (Wald $p < 0.001$), so the Dec coefficients cannot be interpreted as causal. The pre-/post-Act-32 enforcement-break split appears in the Robustness section.

6.1 Pooled-Increase Benchmark

I first report the pooled-Increase specification from Equation (3), which collapses each rate-setter to a single direction tag for rate increases. The pooled coefficients are useful because they approximate the estimand in a conventional local-tax design where the average per-pp dose-response of a municipality-rate change and an SD-rate change without distinguishing whether the rate change moves the worker's bill.

Table 4 reports the per-pp pooled coefficients at the municipality level. The table reports the two coefficients of interest per outcome: the municipality resident-rate Inc effect β^{M}_{Inc} and the school-district-rate Inc analogue β^{SD}_{Inc}.

The muni-rate Inc coefficient on log live-and-work is -0.228^{***}, and the count outcomes move together. Log in-commuters is -0.209^{***} and log total jobs -0.233^{***}, and the share of workers who both live and work in their own municipality is a small -0.019^{**}. The cell decomposition in Section 6.2 shows this aggregate is driven by Inc-bind (-0.573^{***} on log live-and-work) with Inc-becomes and Inc-nobind averaged in. The pooled-Inc ES on log

Table 4: Pooled Inc/Dec headline — municipality level (per pp)

Outcome	Log(work in own muni)	Log(in-commuters)	Log(total jobs)	Share live in muni
Municipality rate (τ^M)				
$\Delta\tau > 0$ (Inc)	-0.228***	-0.209***	-0.233***	-0.019**
	(0.058)	(0.056)	(0.056)	(0.008)
School district rate (τ^{SD})				
$\Delta\tau > 0$ (Inc)	0.128*	0.105	0.115*	0.007
	(0.070)	(0.067)	(0.066)	(0.005)
N	218,677	239,023	239,033	239,033
R^2	0.964	0.959	0.961	0.823
Within R^2	0.001	0.001	0.001	0.001
Dep. mean	3.677	5.847	5.948	0.093
Dep. SD	1.934	1.887	1.901	0.076

Notes: Cluster-robust standard errors in parentheses. Each regression includes both rate-setters' direction-pooled Inc treatments jointly; the reported coefficients are β^{M}_{Inc} and β^{SD}_{Inc}. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

live-and-work fails parallel trends (Wald $p = 0.002$, Table 5), because the Inc-becomes and Inc-nobind munis enter with non-zero pre-event trajectories. The Inc-bind cell on its own passes pre-trends cleanly ($p = 0.33$, Section 6.2.1), which is why the decomposition rather than the pooled coefficient carries the causal claim. The SD-rate Inc coefficients are positive and weak (+0.128* on log live-and-work), sign-reversed relative to the muni-rate finding. Aggregating across a school district mixes the rate-changing municipality (where residents may move out) with non-changing munis in the same SD (which may gain residents from within-SD movers).

6.1.1 Realized-dose ES benchmark diagnostic

Table 5 reports the ATT summary of the muni-level pooled realized-dose event study, with the cluster-robust Wald pre-trend p-value for the joint null $\beta_k = 0$ at $k \leq -2$. Figure 6 plots the event-time path.

The pooled-Inc post-event coefficient on log live-and-work is -0.257, with a small but non-zero pre-event lead, and the joint Wald $p = 0.002$ rejects strong parallel trends. As above, the rejection reflects Inc-becomes and Inc-nobind selection rather than a problem with Inc-bind on its own, which the decomposition isolates.

The other three pooled outcomes in Table 5 and Figure 6 move with the live-and-work

Table 5: Pooled realized-dose event study — municipality level.

Direction	Log(work in own muni)	Log(in-commuters)	Log(total jobs)	Share live in muni
Inc (post)	-0.257	-0.247	-0.271	-0.019
	(0.038)	(0.037)	(0.037)	(0.005)
Inc (pre)	0.052	0.055	0.055	-0.002
	(0.027)	(0.053)	(0.050)	(0.007)
Inc Wald p	0.0021	0.0160	0.0393	0.1424
Diagnostics				
N	163,283	178,335	178,343	178,343
R^2	0.9637	0.9601	0.9621	0.8283
Within R^2	0.0011	0.0015	0.0016	0.0014
Dep. mean	3.6660	5.8580	5.9580	0.0920
Dep. SD	1.9340	1.8800	1.8940	0.0750

Notes: Cluster-robust standard errors in parentheses. Post = mean of β_k for $k = 0, \ldots, 5$; pre = mean for $k \leq -2$. Wald p is the cluster-robust joint test of $H_0\colon \beta_k = 0$ across pre-event k. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

result. The pooled-Inc post on log in-commuters is -0.247 ($p = 0.016$) and on log total jobs -0.271 ($p = 0.039$). The three count measures contract together by roughly a quarter of a log point per pp (0.247–0.271), which points to a single workplace-side margin rather than an outcome-specific artifact. The share of workers who both live and work in the municipality is small and negative (-0.019, $p = 0.14$) and is the only pooled outcome whose joint pre-trend test does not reject. Muni-level share outcomes mix denominator dynamics with binding-channel sorting, and enter as supporting evidence rather than as primary identification.

These pooled magnitudes average over municipalities whose rate changes do and do not move the worker's bill, and three of the four reject the joint pre-trend test. The realized-dose decomposition in Section 6.2 addresses both problems at once. It isolates the Inc-bind cell, where the rate increase actually pushes the home stack above the work rate, from the Inc-becomes and Inc-nobind municipalities that contribute the non-zero leads. There the live-and-work effect sharpens to -0.573^{***} on a cell with clean realized-dose pre-trends (joint Wald $p = 0.33$), while the within-design placebo (Inc-nobind) is null. A large pre-trend-clean effect in the binding cell beside a null placebo is why the decomposition rather than the pooled coefficient carries the causal claim.

Figure 6: Municipality-level pooled realized-dose event study.

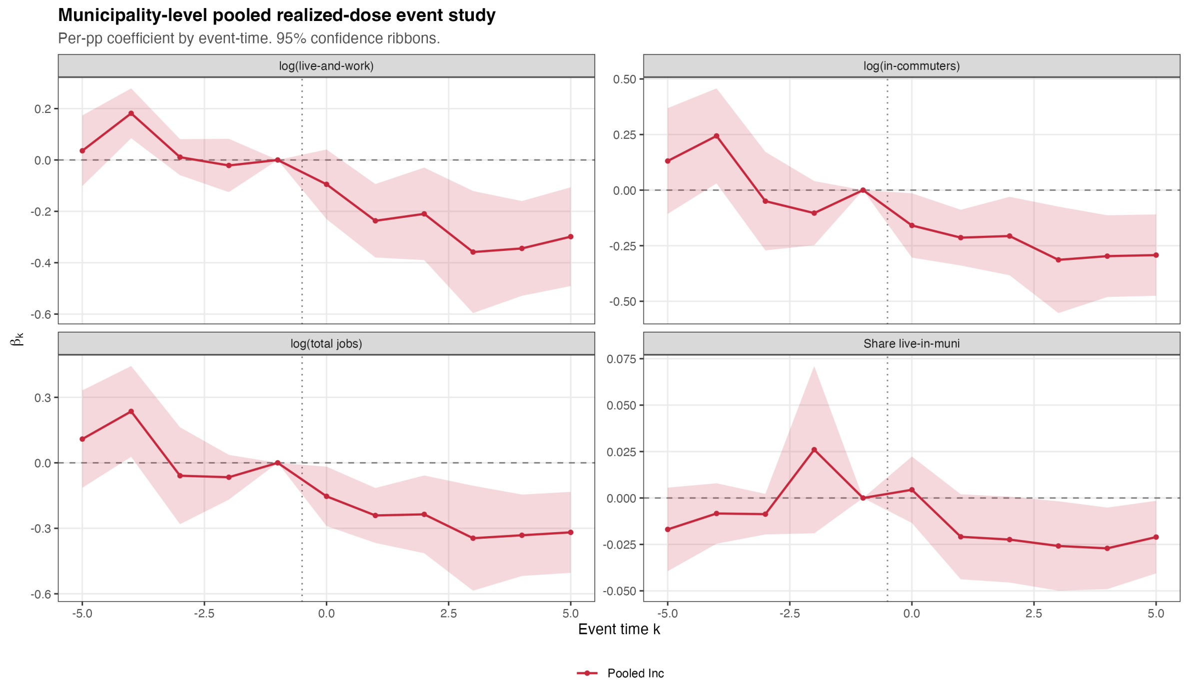

Notes: Per-pp realized-dose coefficient at each event-time for the pooled Inc cohort. Ribbon is the 95% CI; $k = -1$ omitted as the reference period.

6.2 Binding Decomposition

The decomposition is estimated at the municipality level. Both rate-setter treatments enter the same joint regression: each muni-year is assigned its own resident rate τ^M at local level and $\max(\tau^{SD})$ across the school districts the muni reports to. Table 6 reports both sets of cell coefficients from the same joint regression in a single panel: the muni resident-rate cells (top panel) are the headline causal claim, and the SD resident-rate cells (bottom panel) are reported alongside as co-treatment from the identical specification.

The Inc-bind cell has clean realized-dose pre-trends on the three count outcomes (table 7). To understand the results, I translate them using the mean cumulative resident-rate increase at $k = +5$ for the 24 Inc-bind municipalities which is 0.24 pp (Figure 4). At that typical dose, the Inc-bind coefficient on log live-and-work implies a $0.24 \times (-0.573) = -0.14$ log-point change, or 13% reduction in residents who both live and work in their own municipality. The

Table 6: Municipality-level dose-response: muni and SD resident-rate cells (joint regression)

Cell	Log(work in own muni)	Log(in-commuters)	Log(total jobs)	Share live in muni
Muni resident rate (τ^M)				
Inc-bind	-0.573***	-0.412***	-0.490***	-0.058***
	(0.162)	(0.126)	(0.133)	(0.020)
Inc-becomes	-0.250***	-0.247***	-0.266***	-0.014
	(0.056)	(0.055)	(0.056)	(0.011)
Inc-nobind	-0.073	0.233	0.193	-0.037
	(0.143)	(0.302)	(0.268)	(0.032)
School district rate (τ^{SD})				
Inc-bind	0.165	0.159	0.165*	0.005
	(0.109)	(0.096)	(0.093)	(0.007)
Inc-becomes	0.084	0.086	0.096	0.007
	(0.095)	(0.089)	(0.089)	(0.007)
Diagnostics				
N	174,929	191,405	191,413	191,413
R^2	0.963	0.959	0.961	0.825
Within R^2	0.001	0.001	0.001	0.001
Dep. mean	3.659	5.836	5.937	0.092
Dep. SD	1.926	1.883	1.897	0.075

Notes: Stacked-DiD; muni×stack and year×stack FE; SE clustered at muni. Both rate-setters enter the same joint regression; the table reports both sets of cell×dose coefficients. Cells omitted (no estimable coefficients): Inc-nobind. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

dependent-variable mean of 3.66 corresponds to about 39 workers ($e^{3.66}$) in the typical muni, so a 13% drop is roughly 5 fewer live-and-work workers per Inc-bind muni over five years. The same scaling on log in-commuters ($\beta = -0.412$) implies 10% reduction in-commuters, or about 34 fewer in-commuters relative to a typical baseline of 343. The same scaling on log total jobs ($\beta = -0.490$) implies about -11% in jobs located in the municipality; the fact that the estimate for the live-and-work count falls by more than the workplace as a whole (-13% vs -11%) is residential displacement net of the general labor-market contraction.

The share of workers who both live and work in their own municipality is the most direct read on the binding-rule mechanism, and it moves in the predicted direction. The Inc-bind coefficient is -0.058^{***} (Wald $p = 0.003$), so residents' share of local employment falls even net of the labor-market contraction. At the typical Inc-bind event the 0.24 pp cumulative increase implies a $0.24 \times (-0.058)$ approximately a -0.014 change in the share, a drop of about 1.4 percentage points from a baseline mean of 9.2% which is about a 15% relative decline. Against the 380 ($e^{5.94}$) jobs located in a typical municipality that is roughly 5 fewer of those jobs held by residents, and is also the same residential displacement the count outcomes show, now measured net of the overall decline in jobs. I read this effect as causal even though it fails the pre-trends test. This is due to swings in the pre-period that for the most part not statistically significant.

The Inc-becomes cell produces significant coefficients on the same two log count outcomes (-0.250^{***} on log live-and-work; -0.247^{***} on log in-commuters) but the realized-dose pre-trend on both fails the parallel-trends test (Wald $p = 0.0001$ and $p < 0.0001$), though the joint rejection is driven almost entirely by a single anomalous pre-event lead ($k = -4$) rather than a sustained differential trend. I therefore read Inc-becomes as suggestive rather than purely causal.

The Inc-nobind cell is the within-design placebo and has 14 municipalities whose home stack stays below the work rate after the rate hike, so the mechanism predicts that there should not be any movers given the worker's bill never moves. The placebo passes on all three count outcomes: Inc-nobind on log live-and-work is -0.073, on log in-commuters $+0.233$, and on log total jobs $+0.193$, all statistically indistinguishable from zero, so the within-design placebo holds across the count outcomes.

The SD-rate cells in the same regression produce postive effects that are not statistically signficant. The results are thus sign-reversed relative to the muni-rate finding and loses power relative to the aggregate results in table 5. The estimates suggest that in general tax increases from school districts produce positive sorting into the city but cannot decompose that effect for whether it is binding or not. This may partly be due to the SD-rate entering the regression as a co-treatment but does not by itself identify the residential sorting response.

6.2.1 Realized-dose ES decomposition diagnostics

The realized-dose event study disaggregates the pooled pre-trend by cell. Table 7 reports each cell's post-event ATT, pre-event lead, and the joint pre-trend Wald p across the four muni outcomes, and the pre-trend numbers quoted in Section 6.2 are drawn from it. Inc-bind on log live-and-work passes the strong parallel-trends standard. Pre-event coefficients are individually not statistically significant. The post-event path is uniformly negative and deepens through $k = 4$. Inc-becomes and Inc-nobind fail pre-trends ($p < 0.001$ on both). This is what the mechanism predicts. Cohorts selected into rate changes that flip the binding state or stay non-binding are munis on differentially sloped pre-event trajectories. Inc-bind is the only cell that survives the strong parallel-trends standard.

Table 7: Realized-dose event study by binding-state cell — municipality level.

Cell	Log(work in own muni)	Log(in-commuters)	Log(total jobs)	Share live in muni
Inc-bind (post)	-0.762	-0.498	-0.591	-0.069
	(0.095)	(0.080)	(0.076)	(0.015)
Inc-bind (pre)	-0.072	0.119	0.093	-0.020
	(0.068)	(0.088)	(0.080)	(0.012)
Inc-bind Wald p	0.3318	0.6970	0.7925	0.0034
Inc-becomes (post)	-0.171	-0.194	-0.208	-0.012
	(0.037)	(0.039)	(0.038)	(0.004)
Inc-becomes (pre)	0.075	0.081	0.082	-0.002
	(0.028)	(0.057)	(0.054)	(0.008)
Inc-becomes Wald p	0.0001	<0.0001	<0.0001	0.1877
Inc-nobind (post)	0.044	-0.057	-0.068	-0.012
	(0.080)	(0.077)	(0.073)	(0.007)
Inc-nobind (pre)	0.053	-0.782	-0.706	0.069
	(0.151)	(0.244)	(0.216)	(0.027)
Inc-nobind Wald p	<0.0001	<0.0001	<0.0001	<0.0001
Diagnostics				
N	163,014	178,042	178,050	178,050

Notes: Cluster-robust standard errors in parentheses. Post = mean of β_k for $k = 0, \ldots, 5$; pre = mean for $k \leq -2$. Wald p is the cluster-robust joint test of $H_0\colon \beta_k = 0$ across pre-event k. Per-cell decomposition of the muni-level realized-dose ES; each cell is run as a separate regression on {NC, cell}. Inc-bind is the load-bearing cell, with clean pre-trends on the three count outcomes. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 7: Municipality-level realized-dose event study on log live-and-work, by cell.

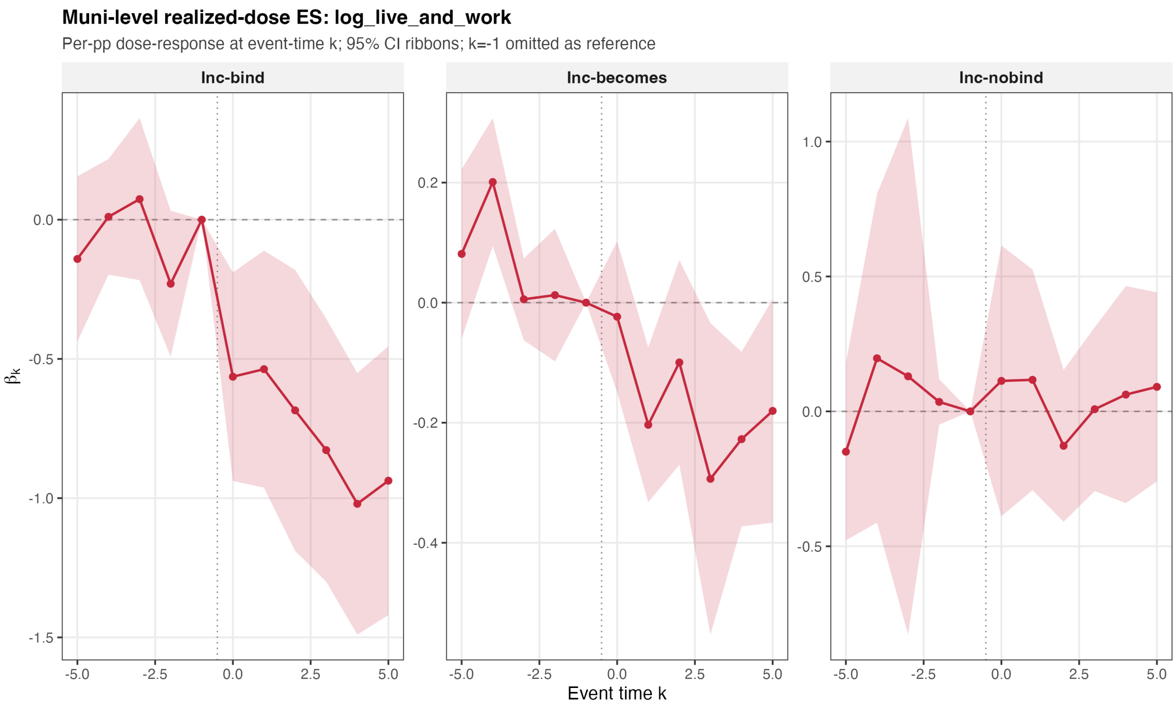

Notes: Per-pp coefficient at each event-time, faceted by cell. Inc-bind pre-event Wald $p = 0.33$; post path uniformly negative, deepening through $k = 4$.

6.3 Summary

The headline finding is from the muni-level binding decomposition: a 1 pp Inc-bind resident-rate increase reduces log live-and-work by -0.573^{***} with a clean realized-dose pre-trend ($p = 0.33$) and a null Inc-nobind within-design placebo (-0.073). The same scaling on log in-commuters gives -0.412^{***}. The pooled Inc benchmark in Section 6.1 averages binding and non-binding events, attenuating the Inc-bind effect to -0.228^{***}, and the decomposition recovers the binding-channel response. Full coefficient tables with fiscal controls are in the appendix.

7 Robustness

This section summarizes the robustness checks behind the headline muni-level dose-response. The realized-dose event-study pre-trend diagnostic was discussed alongside the headline in Section 6.2.1. Full tables behind the robustness summaries are in Appendix C.

7.1 Pre/Post Act 32 Decomposition

Splitting the headline cell-$\times$-dose specification at the Act-32 enforcement break (cohort first event year < 2012 vs ≥ 2012) shows the municipality-level binding-cell effect strengthens after consolidated employer withholding takes effect (-0.593^{**} pre vs -0.546^{***} post on log residents working in own municipality) (see Appendix C.1).

7.2 Act 47 Distressed-Municipality Exclusion

The Inc-bind cohort includes Act 47 distressed cities (Pennsylvania General Assembly, 1987). Act 47 is the legal designation that lets a municipality exceed the 0.5% statutory municipality-EIT cap meaning there may be a selection confounder. An Act 47 designation may come from a period of population decline, (rate hike occurs after the municipality is already losing residents), debt restructuring, and service cuts. Only 2 of the 24 Inc-bind municipalities match the DCED Act 47 designation list (Carbondale and Hazleton). The other historically Act-47-designated PA cities (e.g. Braddock, Coatesville, Johnstown) do not appear in the Inc-bind cell here because their rate trajectories are non-monotone (did not move in one direction) and were removed (Section 5.2). Table 8 re-estimates the municipality-level Inc-bind coefficient on the remaining 22 non-Act-47 Inc-bind municipalities.

Table 8: Act 47 robustness: dropping distressed Inc-bind municipalities from the cousub headline

Outcome	All 24 Inc-bind	Drop 2 Act 47
log live-and-work	-0.718^{***} (0.214)	-0.482^{**} (0.191)
log in-commuters	-0.540^{***} (0.165)	-0.419^{**} (0.172)
log total jobs	-0.621^{***} (0.172)	-0.462^{***} (0.167)
live-in-muni share	-0.059^{***} (0.020)	-0.034^{**} (0.015)

Notes: Muni-level Inc-bind coefficient on each residence-side outcome, with same FE / cluster / controls as the headline. Column 1 is the headline (all 24 Inc-bind municipalities). Column 2 drops 2 Act 47 distressed-designation municipalities from the Inc-bind cohort; the remaining 22 Inc-bind municipalities are mostly Home Rule and Boroughs. Cluster-robust SEs in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The point estimates are direction-preserving and modestly smaller when the two Act 47 munis are dropped. The live-and-work coefficient moves from -0.72 to -0.48, in-commuters from -0.54 to -0.42, total jobs from -0.62 to -0.46, and the live-in-muni share from -0.059 to -0.034. All four coefficients remain statistically significant at 5% or better on the trimmed sample, so the headline is not purely driven by Act 47 selection. The attenuation when Act 47 cities are removed is consistent with those two cities having larger doses.

8 Discussion

The main finding of this paper is that sorting is identified through the binding mechanism. The Inc-bind realized dose of 0.24 pp cumulative over five years, the municipal resident-rate increase reduces the log count of workers who both live and work in the municipality by roughly 13%. The Inc-nobind placebo is not statistically significant, and Inc-becomes lies between the two with a smaller statistically significant effects. The pattern matches the mechanism prediction. The same statutory rate change can produce a strong, moderate, or null behavioral response depending solely on whether the rate moves the worker's bill. This is direct evidence that workers respond to taxes they actually pay rather than to statutory rates per se. The per-pp coefficient (-0.573, or about 44% per pp) is the regression slope and should not be read as a typical-event magnitude. A 1 pp shock is quite far above the median realized dose in this cohort.

Identification at the municipality level is supported by the realized-dose pre-trend diagnostic. The joint Wald test on the Inc-bind cell does not reject parallel trends ($p = 0.33$). The 24-municipality Inc-bind cohort identifies a causal effect, but generally eastern-PA-concentrated set of Home Rule boroughs and Act 47 designees.

Other outcome variables suggest a general decline in economies of these cites such as the in-commuters coefficient is also negative (-0.412^{***}). If residents only relocated their home while keeping their job in the rate-raising municipalities, in-commuters would theoretically rise as someone has to fill the displaced residents' jobs. If Inc-bind munis were raising the work rate $\tau^{\text{work,M}}$ at the same time as the resident rate, in-commuters would mechanically face a higher withholding bill on their own paychecks and the negative coefficient would have a direct work-side explanation. The data does not support that explanation. At the first-event year municipalities raise only the resident rate and leave the work rate unchanged. The negative in-commuters coefficient therefore reflects a broader local labor-market response to municipal fiscal stress, when the binding home-side rate rises rather than a mechanical tax on in-commuters' bills.

The models also incorporate the two-rate-setting mechanisms (municipality and school distircts). The municipality resident-rate Inc-bind coefficient is large and negative, while the SD resident-rate coefficient in the same joint regression is small and positive, but not statistically significant. In the aggregate model the SD rate increases are associated with increased sorting within those cities. Although the models are mixed on statistical significance, the evidence is suggestive of the tax-amenity-bundle model. SD revenues primarily fund schools, which may be viewed as high quality schools thus capitalizing into the willingness to pay to reside in the district. The municipal resident-rate may be less salient to workers as the revenue funds streets, debt service, intergovernmental transfers, and general administration. The sign contrast is suggestive evidence of bundle-level Tiebout sorting and strengthens the binding-channel interpretation rather than complicating it. The binding rule tells us which rate events move workers' bills, and the rate-setter contrast tells us how workers value the revenue once collected.

Six limitations merit acknowledgment. First, the LODES OD data covers PA municipalities but not cross-state inflows in the intra-state main file used here. Incorporating this data would tighten the estimates of in-commuters for border municipalities.

Second, the paper estimates the residential-sorting response to the home-side resident rate but does not separately estimate the work-side response to the municipality work rate ($\tau^{\text{work,muni}}$). The work rate enters the design only through the binding classification of cells. A complementary work-side specification—estimating in-commuter and total-jobs responses to $\tau^{\text{work,muni}}$ events at the municipality level, with a parallel binding decomposition, testing the work side of the mechanism. I defer this to future work.

Third, the per-pp coefficient is a displacement elasticity. The Inc-bind cohort is concentrated in eastern PA (Figure 3). A worker leaving municipality (A) for an adjacent municipality (B) shows up as a -1 in municipality (A) live-and-work count and a $+1$ to that neighbor's residence count. The headline specification partially absorbs neighbor-rate variation through never-changed-only controls and municipality-by-stack fixed effects, but it does not formally model spatial spillovers, and a worker moving to an adjacent municipality whose rate is also moving violates SUTVA. The Butts (2023) estimator would permit me to generalize the stacked DiD to allow for SUTVA violations from spatial spillovers across treated units, using an explicit spatial weighting matrix. Applying this is left for future work.

Fourth, LODES is a stock dataset rather than a person-linked flow dataset. For each (residence municipality, workplace municipality, year), I observe a count of workers, not individual migration flows. The headline coefficient on the municipality count of residents is therefore a net elasticity as measured by the difference of gross outflows and inflows. The paper's evidence that municipal rate hikes reduce the residential count is consistent both with workers leaving and with workers not arriving to replace those who left, however, LODES alone cannot decompose the two. Decomposing into gross outflow and gross inflow would require person-linked administrative data: the IRS Statistics of Income individual master file, the American Community Survey microdata with full residence-history linkage, or the LEHD Job-to-Job Flows file, which tracks the same workers across employer transitions at quarterly frequency. All three are restricted-access products requiring approved-project status under the Census Bureau's Federal Statistical Research Data Centers (FSRDCs) or IRS Statistics of Income agreements. The microdata cannot be exported and must be analyzed at a physical FSRDC location. Replicating the present analysis on person-linked data is a clear extension but not currently being undertaken at this time.

Sixth, the paper does not estimate a comparable elasticity for property taxes. The Oates (1969) capitalization theory predicts a substantially smaller residential-sorting elasticity for property tax than for the EIT, because property tax falls on owners and capitalizes into house prices rather than on workers directly. Estimating a comparable property-tax elasticity would test the falsification implication "is the EIT effect specifically about the EIT, or any local tax raising?" but requires a 2002–2019 municipality-level millage panel that DCED does not publish (the DCED Annual Financial Report schema dropped per-municipality millage in 2006). Constructing this panel from county assessor records is left for future work.

9 Conclusion

This paper answers the question, when do local income taxes cause residential sorting? In Pennsylvania, when the piecewise max-rule makes a statutory rate increase it can mechanically change the worker's tax bill. Under the Inc-bind cell, the increase reduces the count of workers who both live and work in their own municipality by roughly 13% with clean realized-dose pre-trends, and a null finding in the Inc-nobind placebo. This is a local effect for a small cohort of binding rate-raisers rather than a state-average elasticity.

The design delivers the within-design heterogeneity that Agrawal (2026) finds missing from the US local-income-tax literature. The piecewise rule splits otherwise-identical rate increases into cells that moves the worker's bill (Inc-bind) and a placebo cell that does not (Inc-nobind), making the design self-falsifying in a way that pooled local-tax regressions cannot. The municipality and schoold district estimations are opposite with a negative results on the municipality rate and positive on the SD rate. This is consistent with workers sorting on the bundle of taxes and amenities rather than on rates alone.

The natural next step is a companion paper on the work-side response to the muni work rate $\tau^{\text{work,muni}}$, which enters the present design only through the binding classification of cells. More broadly, the framework is portable to any setting where a piecewise tax rule determines whether a statutory change reaches the taxpayer's bill, such as local income taxes and property taxes.

References

Agrawal, D. R. (2015). The tax gradient: Spatial aspects of fiscal competition. *American Economic Journal: Economic Policy*, *7*(2), 1–29. https://doi.org/10.1257/pol.20120360

Agrawal, D. R. (2026). Lessons on State and Local Income Taxes from the Twenty-first Century and Challenges for the Future. *CESifo Working Papers*.

Agrawal, D. R., & Foremny, D. (2019). Relocation of the Rich: Migration in Response to Top Tax Rate Changes from Spanish Reforms. *The Review of Economics and Statistics*, *101*(2), 214–232. https://doi.org/10.1162/rest_a_00764

Bagchi, S. (2022). The Effects of Introducing Withholding on Tax Compliance: Evidence from Pennsylvania's Local Earned Income Tax. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4291598

Baker, A., Callaway, B., Cunningham, S., Goodman-Bacon, A., & Sant'Anna, P. H. C. (2025). Difference-in-Differences Designs: A Practitioner's Guide. https://doi.org/10.48550/ARXIV.2503.13323

Butts, K. (2023). Difference-in-Differences Estimation with Spatial Spillovers. https://doi.org/10.48550/arXiv.2105.03737

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001

de Chaisemartin, C., D'Haultfœuille, X., & Vazquez-Bare, G. (2024). Difference-in-difference estimators with continuous treatments and no stayers. *AEA Papers and Proceedings*, *114*, 393–397. https://doi.org/10.1257/pandp.20241049

Graham, M. R., Kutzbach, M. J., & McKenzie, B. (2014). DESIGN COMPARISON OF LODES AND ACS COMMUTING DATA PRODUCTS.

Pennsylvania Department of Community and Economic Development. (2024a). MunStats portal: Municipal Statistics and Local Government Data. https://munstats.pa.gov/

Pennsylvania Department of Community and Economic Development. (2024b). Statewide Municipal Annual Financial Report (AFR) Database. https://munstats.pa.gov/Reports/ReportInformation2.aspx?report=mAfrForm

Pennsylvania General Assembly. (1932). Sterling Act [Act of Aug. 5, 1932, Sp. Sess., P.L. 45, No. 45]. https://www.legis.state.pa.us/WU01/LI/LI/US/PDF/1932/1/0045..PDF

Pennsylvania General Assembly. (1965). Local Tax Enabling Act, Act 511 [Act of Dec. 31, 1965, P.L. 1257, No. 511. 53 P.S. §§ 6924.101–6924.901]. https://www.legis.state.pa.us/WU01/LI/LI/US/HTM/1965/0/0511..HTM

Pennsylvania General Assembly. (1972). Home Rule Charter and Optional Plans Law [Act of Apr. 13, 1972, P.L. 184, No. 62. 53 Pa.C.S. §§ 2901–3171]. https://www.legis.state.pa.us/WU01/LI/LI/CT/HTM/53/00.029..HTM

Pennsylvania General Assembly. (1987). Municipalities Financial Recovery Act (Act 47) [Act of Jul. 10, 1987, P.L. 246, No. 47. 53 P.S. §§ 11701.101 et seq.]. https://www.legis.state.pa.us/WU01/LI/LI/US/HTM/1987/0/0047..HTM

Pennsylvania General Assembly. (2006). Taxpayer Relief Act (Act 1, Special Session) [Act of Jun. 27, 2006, Sp. Sess. 1, P.L. 1873, No. 1. 53 P.S. §§ 6926.101 et seq.]. https://www.legis.state.pa.us/WU01/LI/LI/US/HTM/2006/1/0001..HTM

Pennsylvania General Assembly. (2008). Act 32 of 2008 (Local Tax Enabling Act — Omnibus Amendments) [Act of Jul. 2, 2008, P.L. 197, No. 32. Amends the Local Tax Enabling Act, 53 P.S. § 6924]. https://www.legis.state.pa.us/WU01/LI/LI/US/HTM/2008/0/0032..HTM

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*, *64*(5), 416–424. https://doi.org/10.1086/257839

U.S. Census Bureau. (2024a). Annual Survey of School System Finances (F-33).

U.S. Census Bureau. (2024b). LEHD Origin-Destination Employment Statistics (LODES), Version 8.

Yang, L., & Heim, B. T. (2017). RESPONSIVENESS OF INCOME TO LOCAL INCOME TAXES: EVIDENCE FROM INDIANA. *National Tax Journal*, *70*(2), 367–391. https://doi.org/10.17310/ntj.2017.2.05

Young, C., Varner, C., Lurie, I. Z., & Prisinzano, R. (2016). Millionaire migration and taxation of the elite: Evidence from administrative data. *American Sociological Review*, *81*(3), 421–446. https://doi.org/10.1177/0003122416639625

A Full Results with Control Coefficients

This appendix reports the muni-level binding-decomposition regression with the fiscal-control coefficients and diagnostics. Table 9 reports the muni-rate and SD-rate cells (two panels) with the shared fiscal controls from the same joint regression.

Table 9: Municipality-level dose-response: muni and SD resident-rate cells with fiscal controls (full spec)

Cell	Log(work in own muni)	Log(in-commuters)	Log(total jobs)	Share live in muni
Muni resident rate (τ^M)				
Inc-bind	-0.573***	-0.412***	-0.490***	-0.058***
	(0.162)	(0.126)	(0.133)	(0.020)
Inc-becomes	-0.250***	-0.247***	-0.266***	-0.014
	(0.056)	(0.055)	(0.056)	(0.011)
Inc-nobind	-0.073	0.233	0.193	-0.037
	(0.143)	(0.302)	(0.268)	(0.032)
School district rate (τ^{SD})				
Inc-bind	0.165	0.159	0.165*	0.005
	(0.109)	(0.096)	(0.093)	(0.007)
Inc-becomes	0.084	0.086	0.096	0.007
	(0.095)	(0.089)	(0.089)	(0.007)
Fiscal controls				
police-services share	0.111	-0.041	-0.031	0.009
	(0.104)	(0.107)	(0.103)	(0.007)
streets share	0.050	0.011	0.014	0.002
	(0.041)	(0.041)	(0.040)	(0.003)
debt-service share	-0.065	-0.029	-0.022	0.004
	(0.054)	(0.060)	(0.058)	(0.005)
EIT-revenue share	-0.036	0.236***	0.220***	-0.013*
	(0.088)	(0.083)	(0.082)	(0.007)
real-estate-tax share	-0.189**	-0.274***	-0.282***	-0.007
	(0.090)	(0.096)	(0.095)	(0.007)
intergovernmental-revenue share	0.018	0.033	0.038	0.005
	(0.046)	(0.050)	(0.050)	(0.004)
Diagnostics				
N	174,929	191,405	191,413	191,413
R^2	0.963	0.959	0.961	0.825
Within R^2	0.001	0.001	0.001	0.001
Dep. mean	3.659	5.836	5.937	0.092
Dep. SD	1.926	1.883	1.897	0.075

Notes: Stacked-DiD; muni×stack and year×stack FE; SE clustered at muni. Both rate-setters enter the same joint regression; the table reports both sets of cell×dose coefficients with the shared fiscal controls. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B Realized-Dose Event Studies, Full Set

Per-cell realized-dose Sun–Abraham event studies at the muni level for the remaining residence-side outcomes (log live-and-work is in the body, §6.2.1).

Figure 8: Muni-level realized-dose ES on log(in-commuters), by cell.

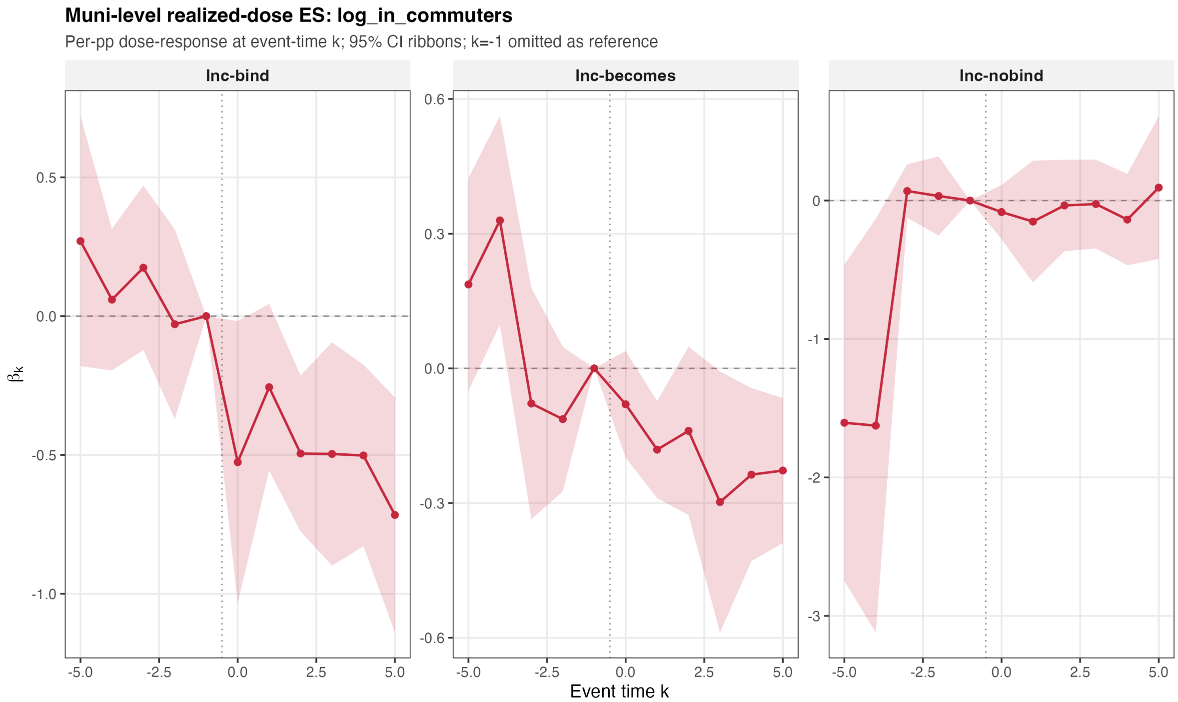

Figure 9: Muni-level realized-dose ES on log(total jobs), by cell.

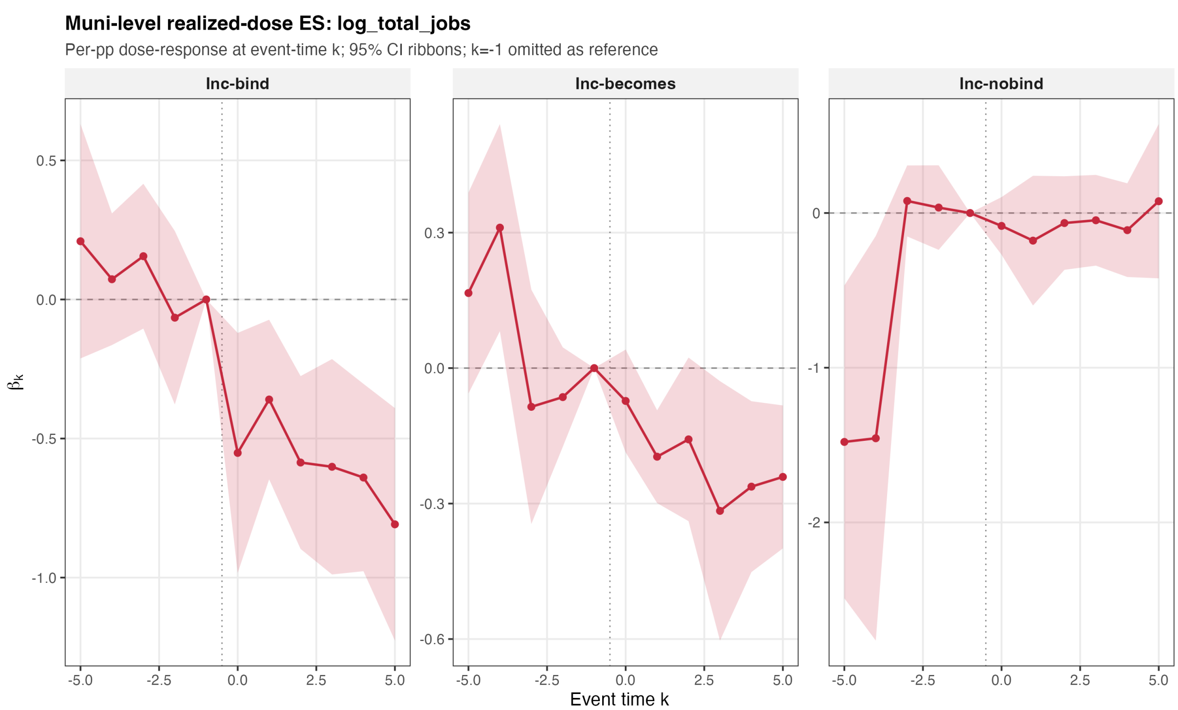

Figure 10: Muni-level realized-dose ES on the share of workers living in their own municipality, by cell.

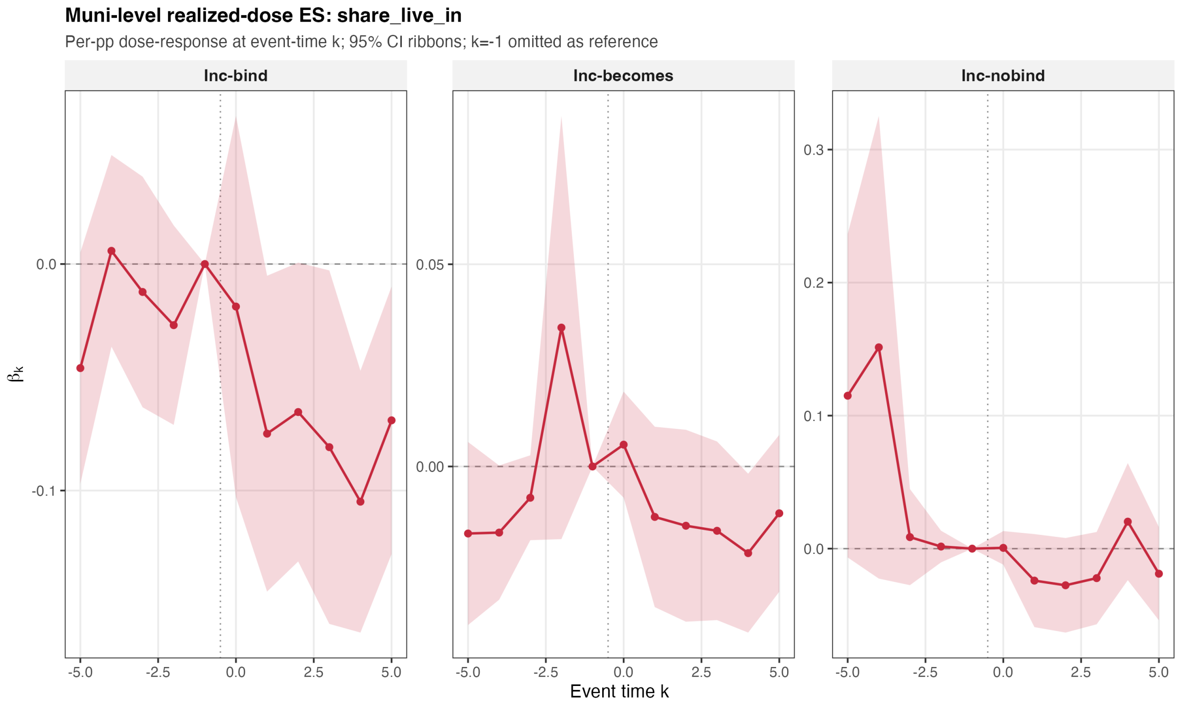

C Robustness: Full Tables

Full tables behind the robustness summaries in Section 7.

C.1 Pre/post Act 32 split

Headline cell × dose specification re-estimated on pre-Act-32 ($c < 2012$) and post-Act-32 ($c \geq 2012$) cohort subsamples at the municipality level.

Table 10: Pre-/post-Act-32 split — municipality level

Outcome	Log(work in own muni)	Log(in-commuters)	Log(total jobs)	Share live in muni
Inc-bind				
Pre-Act 32	-0.593**	-0.404**	-0.500***	-0.068**
	(0.235)	(0.180)	(0.189)	(0.030)
Post-Act 32	-0.546***	-0.430***	-0.481***	-0.042***
	(0.168)	(0.147)	(0.145)	(0.005)
Inc-becomes				
Pre-Act 32	-0.319***	-0.286**	-0.343***	-0.042*
	(0.121)	(0.116)	(0.126)	(0.023)
Post-Act 32	-0.208***	-0.222***	-0.215***	0.005
	(0.055)	(0.050)	(0.047)	(0.009)
Inc-nobind				
Post-Act 32	-0.074	0.231	0.191	-0.037
	(0.143)	(0.302)	(0.267)	(0.032)
Pre-Act 32				
N	174,776	191,249	191,257	191,257
R^2	0.963	0.959	0.961	0.825
Within R^2	0.001	0.001	0.001	0.001
Dep. mean	3.658	5.835	5.935	0.092
Dep. SD	1.925	1.882	1.896	0.075
Post-Act 32				
N	174,614	191,069	191,077	191,077
R^2	0.963	0.959	0.961	0.825
Within R^2	0.001	0.001	0.001	0.001
Dep. mean	3.656	5.834	5.934	0.092
Dep. SD	1.925	1.881	1.896	0.075

Notes: Cluster-robust standard errors in parentheses. Treated cohorts split at first event year $c < 2012$ (Pre-Act 32) vs $c \geq 2012$ (Post-Act 32). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.