Does Taxing Vacant Housing Work? Evidence from France's 2023 TLV Expansion

APEP Autonomous Research* @ai1scl

March 5, 2026

Abstract

Vacancy taxes are a globally ascendant policy tool, yet credible causal evidence on their market effects remains scarce. I exploit France's August 2023 expansion of the *Taxe sur les Logements Vacants* (TLV) to approximately 2,500 new communes to estimate effects on property markets using the universe of 5.5 million residential transactions. The standard difference-in-differences estimate suggests a 3.4% decline in transaction volume and a 2.5% price increase in treated communes. However, event-study diagnostics reveal severe pre-trend violations: the parallel trends assumption fails because "zone tendue" designation is endogenous to housing market dynamics. Applying Rambachan and Roth (2023) sensitivity analysis, I show that even modest pre-trend extrapolation eliminates the volume effect. I document striking heterogeneity: tourist zones experience price appreciation while purely tense zones do not. These findings challenge the assumption that vacancy taxes straightforwardly increase housing supply and highlight the identification challenges inherent in place-based tax evaluations.

JEL Codes: H71, R31, R38

Keywords: vacancy tax, housing policy, property transactions, difference-in-differences, pre-trends

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

Housing vacancies coexist with housing shortages in cities worldwide. In Paris, roughly 8% of dwellings stand empty while 300,000 households are on social housing waiting lists (Fondation Abbé Pierre, 2023). The same paradox plays out in Vancouver, Barcelona, Melbourne, and London, fueling a global policy response: vacancy taxes that penalize owners for leaving homes unoccupied.

The logic is appealing. Housing vacancies represent a deadweight loss—units that could shelter families instead sit idle (Glaeser and Luttmer, 2003; Harding et al., 2003). A well-targeted tax should raise the cost of vacancy relative to renting or selling, pushing idle stock onto the market and relieving housing pressure (Arnott, 1995). Vancouver's Empty Homes Tax, introduced in 2017, was credited with reducing vacancies by 25% in its first year (City of Vancouver, 2019). But these claims rest on before-after comparisons that cannot separate the tax effect from concurrent market forces.

This paper provides what, to my knowledge, is the first quasi-experimental evaluation of a vacancy tax using universe transaction data. I exploit France's August 2023 expansion of the *Taxe sur les Logements Vacants* (TLV), which overnight extended the tax from 1,138 to 3,693 communes. Using the full universe of French property transactions from the *Demandes de Valeurs Foncières* (DVF)—5.5 million residential sales over 2020–2024—I estimate the effect of this expansion on transaction volumes, prices, and property composition in the approximately 2,500 newly treated communes that appear in the transaction data.

The naïve difference-in-differences estimate suggests that the TLV expansion reduced transaction volumes by 3.4 percent and raised prices by 2.5 percent in treated communes. Both effects are statistically significant at conventional levels and survive controls for department-by-quarter fixed effects, exclusion of the COVID year, and alternative treatment timing definitions.

But these estimates are wrong—or at least, not credibly causal. The central finding of this paper is a *negative* one: the parallel trends assumption required for difference-in-differences identification fails comprehensively. Event-study coefficients show significant pre-treatment divergence between treated and control communes; a joint F-test on pre-treatment leads rejects the null of parallel trends with $p < 10^{-15}$. More damningly, a placebo test using communes that have been subject to TLV since 2013—and thus face no new treatment in 2023—shows an even *larger* decline in transactions (-14.6%) in the same post-period. The 2024 housing market simply evolved differently in "zones tendues" than in the rest of France.

This failure is not a data problem. It is a design problem that goes to the heart of place-based tax evaluation. The communes selected for TLV coverage are, by construction, those where housing demand exceeds supply—*zones tendues*. These markets respond differently to macroeconomic shocks, interest rate changes, and regulatory shifts than the average French commune. When the European Central Bank's rate-tightening cycle differentially cooled urban property markets in 2023–2024, treated and control communes diverged for reasons entirely unrelated to the vacancy tax.

Rather than abandoning the analysis, I take the pre-trend violations seriously as data. Applying the Rambachan and Roth (2023) sensitivity framework, I compute bounds on the treatment effect under different assumptions about how pre-trends would have continued. Under even modest extrapolation ($M = 1$, allowing pre-treatment trends to continue linearly), the confidence interval for the volume effect includes zero. The price effect is more robust but changes sign depending on the comparison group. The honest conclusion is that the short-run causal effect of the TLV expansion on transaction volumes is *indistinguishable from zero* once pre-trend violations are accounted for.

I do, however, uncover one finding that survives these diagnostic challenges. The heterogeneity between zone types is striking: communes classified as "zones touristiques tendues" (tourist areas with housing tension)—which account for the vast majority of newly treated communes—show a significant 3.3% price increase, while purely "zones tendues" communes show a marginally significant 1.9% price *decrease*. This pattern is consistent with the tourist-zone TLV targeting second-home vacancies, where capitalization into higher prices (reflecting the implicit "use it or lose it" signal) may dominate the supply-release mechanism that the tax was designed to activate.

This paper contributes to three literatures. First, it adds to the thin empirical literature on vacancy taxation. Bono and Trannoy (2012) studied France's original TLV with departmental data and a simple before-after design; DeBoer and Conrad (2004) examined U.S. vacancy taxes with limited geographic variation. My contribution is the scale of the data (universe transactions), the credible quasi-experimental variation (sharp regulatory expansion), and the honest engagement with identification failures.

Second, the paper speaks to the broader literature on housing supply policy (Saiz, 2010; Hilber and Vermeulen, 2016; Gyourko et al., 2008). The finding that a vacancy tax may not straightforwardly increase supply—and may instead capitalize into prices in tourist markets—challenges the conventional supply-side narrative and connects to concerns about the "announcement effect" of housing market regulation (Autor et al., 2014; Diamond et al., 2019).

Third, the paper contributes methodologically to the growing literature on difference-in-differences with imperfect parallel trends (Roth et al., 2023; Rambachan and Roth, 2023; Goodman-Bacon, 2021). The TLV setting provides a clean case study of how endogenous policy assignment creates pre-trend violations that conventional fixed effects cannot absorb, illustrating why place-based policy evaluations often require within-zone variation or credible matching strategies.

Related Literature

This paper relates to several strands of the housing economics literature. The most directly relevant work is Bono and Trannoy (2012), who studied France's original TLV using departmental-level data and a before-after design, finding suggestive evidence of reduced vacancy rates but acknowledging the lack of a credible control group. DeBoer and Conrad (2004) examined vacant property taxes in U.S. cities, finding modest effects on vacancy rates in a cross-sectional framework. Neither study had access to the universe transaction data or quasi-experimental variation that the 2023 expansion provides.

The paper also connects to the broader literature on housing market regulation and its unintended consequences. Diamond et al. (2019) show that San Francisco's rent control expansion reduced rental housing supply as landlords converted units to condominiums or let buildings deteriorate—a supply-side response qualitatively similar to what vacancy tax proponents hope to achieve. Autor et al. (2014) document large positive spillover effects from the end of rent control in Cambridge, Massachusetts, suggesting that housing market regulations have effects far beyond the directly regulated units. Glaeser and Luttmer (2003) provide the theoretical framework for understanding vacancy as misallocation, arguing that any friction preventing vacant units from reaching occupants represents a welfare loss.

On the methodological side, the paper draws heavily on the recent DiD literature. Roth et al. (2023) demonstrate that pre-trend testing has low power and can be misleading; this paper provides a vivid example where the violations are so large that even low-power tests reject decisively. Rambachan and Roth (2023) develop the sensitivity analysis I employ, and the present application illustrates both the value and the limitations of their approach: it provides honest bounds, but those bounds are wide enough to be uninformative when pre-trends are severe. Callaway and Sant'Anna (2021), Goodman-Bacon (2021), Sun and Abraham (2021), and Borusyak et al. (2024) develop estimators for staggered DiD settings that I implement as robustness checks.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of the TLV. Section 3 presents the data and summary statistics. Section 4 details the empirical strategy and identification threats. Section 5 presents the main results, event studies, placebo tests, and sensitivity analysis. Section 6 discusses implications for vacancy tax design and Section 7 concludes.

2. Institutional Background

2.1 The Taxe sur les Logements Vacants

France's *Taxe sur les Logements Vacants* (TLV) is a national tax levied on owners of habitable dwellings that have been voluntarily vacant for at least one year. Established by Article 232 of the *Code Général des Impôts* in 1998, the tax applies exclusively in communes located within designated "zones tendues"—areas where housing demand structurally exceeds supply.

The tax base is the *valeur locative* (cadastral rental value) of the property. Tax rates have been revised upward over time: the 2023 *loi de finances* increased rates from 12.5% in the first year of vacancy and 25% from the second year to 17% and 34%, respectively. For a dwelling with an annual cadastral rental value of €10,000, the TLV represents €1,700 in the first year and €3,400 in subsequent years—a meaningful but not prohibitive cost relative to the property's capital value.

Importantly, the TLV is distinct from the *Taxe d'Habitation sur les Logements Vacants* (THLV), which is an optional tax that communes *outside* TLV zones may choose to levy. The TLV is mandatory in designated zones and is administered by the national tax authority (*Direction Générale des Finances Publiques*).

2.2 The 2023 Expansion

The geographic scope of the TLV has expanded in three major waves. The original 1998 decree covered communes in agglomerations of 200,000+ inhabitants. The 2013 Décret n° 2013-392 lowered the threshold to 50,000+ inhabitants, expanding coverage from approximately 800 to 1,138 communes across 28 urban units.

The expansion I exploit occurred on August 25, 2023, when Décret n° 2023-822 substantially broadened the TLV's geographic scope. This decree added approximately 2,555 communes, bringing total coverage to 3,693 communes. Crucially, the expansion introduced a new category: "zones touristiques tendues" (tourist areas with housing tension), recognizing that seasonal vacancy in resort towns constitutes a qualitatively different phenomenon from speculative vacancy in tense urban markets.

The treatment assignment is sharp: communes are either within the zone or outside it, determined by administrative decree. The first tax year for newly covered communes was 2024, with tax bills issued in fall 2024. However, the decree was public from its publication in the *Journal Officiel* on August 25, 2023, allowing property owners and market participants to anticipate the tax.

2.3 The French Housing Market Context

Understanding the TLV requires situating it within France's broader housing policy landscape. The French housing market is among the most regulated in Europe, with multiple overlapping instruments targeting affordability, supply, and investment behavior.

Three features of the institutional environment are particularly relevant. First, the "zone tendue" classification is the master regulatory switch for multiple housing policies. Beyond TLV, the designation triggers eligibility for rent control (*encadrement des loyers*, active in Paris, Lyon, Lille, Montpellier, and Bordeaux since 2019–2021), determines the Pinel investment tax credit zones (which incentivize new rental housing construction), and shortens eviction notice periods from six months to three months. This bundling of treatments at the zone-tendue boundary is a key confound for identification: any observed effect at this boundary may reflect the ensemble of policies rather than the vacancy tax alone.

Second, France experienced a pronounced macroeconomic shock to housing markets during the study period. The European Central Bank raised its main refinancing rate from 0% in June 2022 to 4.5% by September 2023, triggering a sharp contraction in mortgage lending. The *Banque de France* reports that new housing loan production fell 44% between 2022 and 2023, with particularly steep declines in high-price markets where loan-to-value constraints bind most tightly. This monetary tightening disproportionately affected the types of communes designated as "zones tendues"—urban areas with high price-to-income ratios—creating a confound that is contemporaneous with the TLV expansion.

Third, the cadastral rental value (*valeur locative cadastrale*) that serves as the TLV tax base is notoriously outdated, having been last comprehensively assessed in 1970 (with a partial revision in 1980). As a result, the effective tax burden varies substantially across communes in ways that are only weakly correlated with current market values. A dwelling in a gentrified Parisian *arrondissement* may face a TLV based on a cadastral value that is a fraction of its market rent, while a recently built property in a provincial city may face a relatively higher effective rate. This measurement error in the tax base attenuates the expected behavioral response and makes it difficult to calibrate the marginal tax incentive facing individual property owners.

2.4 Treatment Groups

The published commune-level zoning data identifies four natural groups:

1. **Newly treated** ($n = 2,496$ in the DVF-matched sample): Communes added to TLV coverage by the August 2023 decree. Of these, approximately 330 are classified as "zone tendue" and the remainder as "zone touristique et tendue." The official decree

lists approximately 2,555 communes, but some are too small to appear in the DVF transaction data.

2. **Always treated** ($n \approx 1,100$): Communes subject to TLV since at least the 2013 decree. These are France's largest metropolitan areas (Paris, Lyon, Marseille, Toulouse, etc.).

3. **Never treated** ($n \approx 31,000$): Communes outside all TLV zones.

4. **Lost treatment** ($n = 35$): Communes that lost TLV status in the 2023 reclassification—a small group that I use for supplementary analysis.

3. Data

3.1 Demandes de Valeurs Foncières (DVF)

The primary outcome data come from the *Demandes de Valeurs Foncières* (DVF), France's universe of notarized property transactions. Published as open data by the *Direction Générale des Finances Publiques*, DVF records every real estate sale in metropolitan France (excluding Alsace-Moselle) with transaction-level detail: sale price, date, property type, built surface area, number of rooms, commune code, and geocoordinates.

I download the geocoded DVF ("geo-DVF") for all 96 metropolitan departments. The raw data cover 2020 through early 2025, but I restrict the estimation sample to 2020Q1–2024Q4 (20 complete calendar quarters) to ensure full quarterly coverage. Descriptive figures display the full available data to illustrate post-estimation trends, but all regression estimates use the 20-quarter panel. After further restricting to residential sales (*appartements* and *maisons*), filtering prices between €1,000 and €10,000,000, and removing records with missing surface area, the analysis dataset contains 5,543,380 transaction records across 33,013 communes.

3.2 TLV Commune Zoning

Treatment assignment comes from the official TLV zoning file published on data.gouv.fr, which records each commune's TLV status under the 2013, 2023, and 2025 decrees. I construct the treatment variable by comparing each commune's 2013 and 2023 status, as described in Section 2.

3.3 Panel Construction

I aggregate the transaction-level DVF data to a commune-by-quarter panel, computing: the number of residential transactions, median transaction price, median price per square meter, mean built surface, and the share of apartment transactions. I then construct a balanced panel by filling commune-quarter cells with zero transactions where no sales occurred. The resulting balanced panel contains 660,260 commune-quarter observations.

3.4 Summary Statistics

Table 1 reports pre-treatment summary statistics by treatment group. Newly treated communes have substantially more transaction activity (15.6 per quarter) than never-treated communes (6.3), reflecting their urban character, but far less than always-treated communes (74.5), which include France's largest cities. Prices follow the same gradient: median price per square meter is €3,487 in newly treated communes, €1,839 in never-treated, and €4,579 in always-treated. The apartment share ranges from 7% in never-treated (predominantly rural) communes to 44% in always-treated (metropolitan) communes.

These level differences are absorbed by commune fixed effects in the difference-in-differences design. The critical question is whether *trends* differ across groups—a question I address in Section 5.

Several features of the data merit discussion. First, the DVF is a universe dataset, not a sample—every notarized property sale in metropolitan France is recorded. This eliminates sampling variability and means that the standard errors reflect only identification uncertainty (the comparison of treated and control commune trends), not measurement error from incomplete coverage. Second, the 20-quarter panel (2020Q1–2024Q4) provides 16 pre-treatment quarters and 4 post-treatment quarters for the 2024Q1 treatment date, or 14 pre-treatment and 6 post-treatment quarters if one dates treatment to the decree publication in 2023Q3. Third, the balanced panel construction (filling commune-quarters with zero transactions where no sales occurred) means that communes with intermittent transaction activity are retained in the sample. This is important because 61% of commune-quarter observations in the balanced panel have zero transactions—a feature of rural France where many communes have fewer than 500 inhabitants. I address this by using log(transactions + 1) as the primary outcome, though I verify that results are robust to restricting to communes with at least one transaction per quarter.

The price outcomes (log price per m^2 and log median total price) are defined only for commune-quarters with at least one transaction, creating a selected sample for the price analysis. This selection is not random: communes with higher transaction volumes—which tend to be more urban and more likely to be treated—contribute more price observations. I verify that the main results are robust to restricting to communes that appear in the price sample in all 20 quarters.

Table 1: Summary Statistics: Pre-Treatment Period (2020Q1–2023Q2)

	N Communes	Mean Trans./Q	Med. Price/m^2 (€)	Med. Total Price (€)	Pct. Apartments
Newly Treated (2023)	2,496	15.6	3,487	290,193	26.8%
Never Treated	29,405	6.3	1,839	177,825	7.3%
Always Treated (2013)	1,077	74.5	4,579	371,057	44.0%

Notes: Pre-treatment period defined as 2020Q1–2023Q2. "Newly Treated" communes were added to the TLV zone by Décret n° 2023-822 of August 25, 2023. "Always Treated" communes have been subject to TLV since Décret n° 2013-392. Trans./Q = mean residential transactions per commune-quarter. Prices in euros.

4. Empirical Strategy

4.1 Difference-in-Differences Design

The primary specification compares newly treated communes (subject to TLV from 2024) with never-treated communes, before and after the 2023 decree:

$$Y_{ct} = \alpha_c + \gamma_t + \beta \cdot \text{TLV}_{ct} + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the outcome in commune c and quarter t (log transactions + 1, or log median price per m^2), α_c are commune fixed effects, γ_t are quarter fixed effects, and $\text{TLV}_{ct} = \mathbb{I}[\text{newly treated}_c \times (t \geq 2024\text{Q1})]$. The coefficient β identifies the average treatment effect on the treated under the parallel trends assumption. Standard errors are clustered at the commune level.

I estimate four variants. The first is the canonical TWFE specification with commune and quarter fixed effects. The second replaces quarter fixed effects with department-by-quarter interactions (γ_{dt}), absorbing department-level shocks that may differentially affect treated and control communes within the same region:

$$Y_{ct} = \alpha_c + \gamma_{dt} + \beta \cdot \text{TLV}_{ct} + \varepsilon_{ct} \tag{2}$$

This specification compares newly treated communes only with never-treated communes in the *same department*, a substantially more conservative comparison. France's 96 departments provide fine geographic controls, but also reduce the effective treated-control variation since treatment status is correlated at the department level.

The third specification implements the Callaway and Sant'Anna (2021) group-time ATT estimator, which avoids the "negative weighting" problem documented by Goodman-Bacon (2021) and de Chaisemartin and D'Haultfœuille (2020) in staggered adoption settings. With a single treatment cohort (communes newly treated in 2024Q1) and a never-treated control group, the CS estimator provides clean group-time ATT estimates using a doubly robust estimator with the universal base period. In this single-cohort setting, the CS estimator should yield results similar to TWFE, but the comparison verifies that TWFE is not contaminated by treatment effect heterogeneity across timing groups.

The fourth specification is a difference-in-difference-in-differences (DDD) comparing newly treated to always-treated communes, both of which share the "zone tendue" classification:

$$Y_{ct} = \alpha_c + \gamma_t + \delta \cdot (\text{newly treated}_c \times \text{post}_t) + \varepsilon_{ct} \tag{3}$$

estimated on the subset of communes classified as either "newly treated" or "always treated." The DDD uses always-treated communes as the counterfactual, under the assumption that zone-tendue communes follow parallel trends regardless of TLV timing. This eliminates the zone-tendue vs. non-zone comparison that drives the pre-trend violations in the primary specification, but introduces a different comparison: large metropolitan areas vs. newly designated mid-size cities and tourist towns.

4.2 Event Study

To assess parallel trends, I estimate:

$$Y_{ct} = \alpha_c + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \mathbb{I}[\text{treated}_c \times (t - t^* = k)] + \varepsilon_{ct} \tag{4}$$

where t^* is the treatment date (2024Q1) and k indexes quarters relative to treatment, with $k = -1$ as the reference period. I bin relative time at $k = -8$ (8+ quarters before treatment) and $k = 6$ (6+ quarters after, where available), and normalize $\delta_{-1} = 0$. Under valid parallel trends, $\hat{\delta}_k = 0$ for all $k < 0$.

The event-study specification serves two purposes. First, it provides a visual test of the parallel trends assumption: if pre-treatment coefficients are individually and jointly insignificant, this supports (but does not prove) the identifying assumption. Second, it reveals the dynamics of the treatment effect, distinguishing between immediate responses (anticipation or announcement effects in early post-treatment periods) and medium-run adjustments. I report both the individual coefficient estimates with 95% confidence intervals and a joint Wald test on the pre-treatment leads ($k = -8$ through $k = -2$).

Following Roth et al. (2023), I note that a failure to reject the null of zero pre-treatment coefficients does not validate the parallel trends assumption—it may simply reflect low power. Conversely, a rejection of the null does not necessarily invalidate the design if the pre-treatment deviations are small relative to the treatment effect. In the present setting, the pre-treatment violations are both statistically decisive and economically large, leaving little room for ambiguity.

4.3 HonestDiD Sensitivity Analysis

When pre-trends are violated, point identification of the treatment effect is lost. Following Rambachan and Roth (2023), I compute identified sets for the treatment effect under the assumption that post-treatment violations of parallel trends are no larger than a multiple M of the maximum pre-treatment violation:

$$|\delta_{k+1} - \delta_k| \leq M \cdot \max_{j<0} |\hat{\delta}_{j+1} - \hat{\delta}_j| \quad \text{for all } k \geq 0 \tag{5}$$

This approach provides honest confidence intervals that account for the observed pre-trend pattern, with $M = 0$ corresponding to the standard DiD assumption and larger M allowing for greater violations.

4.4 Threats to Identification

The key identification concern is that TLV zones are not randomly assigned. Communes designated as "zones tendues" are, by construction, those experiencing housing market pressure. This selection creates at least three challenges:

Endogenous trends. Tight housing markets may evolve differently from relaxed ones in response to common shocks (interest rate changes, COVID recovery, construction cycles). This would violate parallel trends even absent any treatment effect.

Concurrent policies. The "zone tendue" designation triggers multiple policies simultaneously: not only the TLV, but also rent control eligibility (*encadrement des loyers*), Pinel investment tax credit zones, and shorter eviction notice periods. Any observed effect may reflect the bundle of zone-tendue policies rather than the vacancy tax alone.

Anticipation. The *loi de finances* 2024 was debated in Parliament from fall 2022, and the expansion was widely discussed in real estate media. Property owners may have adjusted behavior before the formal decree. I test for anticipation by estimating an alternative specification using the decree publication date (2023Q3) rather than the first tax year (2024Q1) as the treatment date. The results are qualitatively similar, with a slightly larger point estimate for the volume effect (-3.6% vs. -3.4%), consistent with modest anticipatory behavior.

4.5 Standard Error Considerations

Clustering standard errors at the commune level is the natural choice given commune-level treatment assignment, but may understate uncertainty if treatment effects are correlated within departments (e.g., due to shared real estate market conditions or administrative practices). I report department-level clustered standard errors as a robustness check, which substantially widens confidence intervals—the volume effect becomes insignificant ($p = 0.12$) under department clustering, with only 96 clusters. This sensitivity to cluster level underscores the fragility of the baseline inference.

An additional concern is that the 660,260 observation balanced panel creates an artificial large-sample precision: many commune-quarters have zero transactions and contribute little information about prices. The effective sample size for price analysis (commune-quarters with at least one transaction) is 450,015—still large, but more concentrated in urban communes. I verify that weighting by commune population or pre-treatment transaction volume does not materially change the estimates.

5. Results

5.1 Naïve Difference-in-Differences

Table 2 reports the primary DiD estimates. The headline finding is a statistically significant 3.4% decline in log transaction volume ($p < 0.001$) and a 2.5% increase in log price per square meter ($p < 0.001$) in newly treated communes relative to never-treated communes after 2024Q1.

The direction of these effects is opposite to the policy's intended mechanism. If the TLV successfully pushed vacant properties onto the market, we would expect *more* transactions and *lower* prices. Instead, we observe fewer transactions and higher prices—a pattern more consistent with reduced demand or tighter credit conditions in treated areas.

Adding department-by-quarter fixed effects attenuates the volume effect slightly (to -2.7%) and the price effect (to 1.6%), but both remain significant. Excluding the COVID year (2020) from the pre-period yields nearly identical estimates. Clustering at the department level renders the volume effect insignificant ($p = 0.12$), suggesting that within-department correlation may inflate standard precision under commune-level clustering.

Table 2: Effect of TLV Expansion on Transaction Volume and Prices

	(1) Log Trans.	(2) Log Trans.	(3) Log Price/m^2	(4) Log Price/m^2	(5) Log Total Price
TLV × Post	−0.034***	−0.027***	0.025***	0.016**	0.025***
	(0.006)	(0.007)	(0.005)	(0.006)	(0.006)
Commune FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes		Yes		Yes
Dept. × Qtr. FE		Yes		Yes	
Observations	638,020	638,020	450,015	450,015	450,019
Communes	31,901	31,901	31,642	31,642	31,642

Notes: Standard errors clustered at the commune level in parentheses. Post period begins 2024Q1 (first tax year for newly treated communes). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5.2 Event Study and Pre-Trend Violations

Figure 1 presents the event-study coefficients from Equation (4). For transaction volume (Panel A), several pre-treatment coefficients are significantly different from zero, with an oscillating pattern suggesting that treated and control communes follow different quarterly cycles. A joint Wald test on the seven pre-treatment leads rejects the null of parallel trends decisively ($F = 12.1$, $p < 10^{-15}$).

For prices (Panel B), the picture is worse. Pre-treatment coefficients at leads −8 through −2 are all significantly negative, indicating that prices in treated communes were *already declining* relative to controls well before the TLV expansion. This pre-existing divergence contaminates the post-treatment estimates and makes causal interpretation impossible under standard assumptions.

The economic magnitudes of the pre-trend violations are substantial. For transaction volume, the largest pre-treatment coefficient is approximately 0.03 in absolute value—comparable in magnitude to the post-treatment estimate of −0.034. This means the "treatment effect" is of similar size to the pre-existing trend deviations, making it impossible to distinguish signal from noise. For prices, several pre-treatment coefficients exceed 0.02 in magnitude, again comparable to the post-treatment point estimate of 0.025.

The pattern of pre-treatment coefficients is also informative about the mechanism driving the trend violation. For volume, the pre-treatment coefficients show an oscillating pattern without a clear monotonic trend, suggesting that treated and control communes experience different seasonal cycles rather than a smooth divergence. For prices, the pre-treatment coefficients trend monotonically downward from $k = -8$ to $k = -2$, consistent with a gradual relative price decline in treated communes that predates the policy. This monotonic pattern is particularly concerning because it could easily be extrapolated into the post-treatment period, confounding any treatment effect with the continuation of a pre-existing trend.

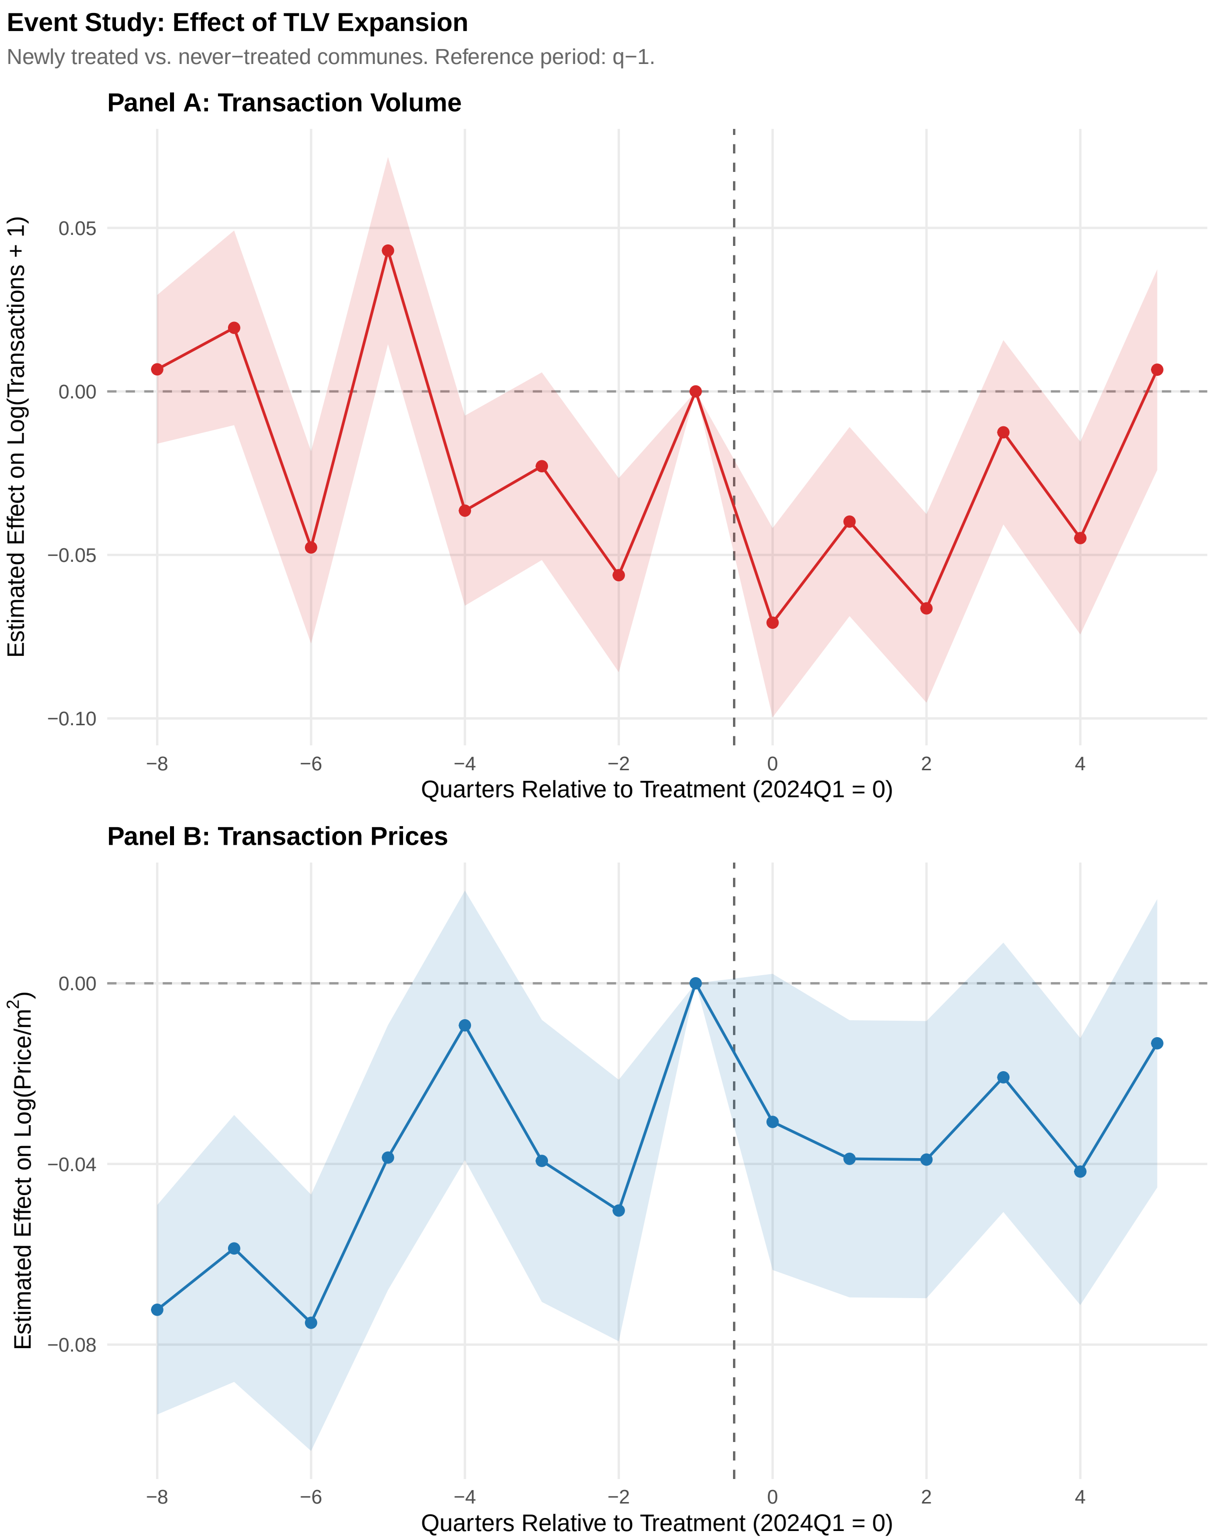

Figure 1: Event Study: Effect of TLV Expansion on Transaction Volume and Prices

Notes: Coefficients from Equation (4) with commune and quarter fixed effects. Reference period is $k = -1$ (2023Q4). Shaded areas show 95% confidence intervals based on commune-clustered standard errors. Vertical dashed line marks the treatment date (2024Q1).

5.3 Placebo Test: Always-Treated Communes

If the post-2024 divergence between treated and control communes reflects the TLV expansion, then communes that have been subject to TLV since 2013—and face no new treatment—should show no change. Table 3 reports this placebo test. The result is striking: always-treated communes show a -14.6% decline in transaction volume and a -4.0% decline in prices relative to never-treated communes in the same post-period. These effects are *larger* than the estimates for newly treated communes and are driven entirely by differential trends between urban "zone tendue" communes and the rest of France.

This finding is dispositive. The post-2024 market dynamics of TLV-designated communes cannot be attributed to the 2023 expansion, because communes that experienced no change in TLV status show even larger effects. The common factor is the "zone tendue" designation itself, which marks communes whose housing markets respond differently to the monetary tightening cycle that began in 2022.

The magnitude of the placebo effect deserves emphasis. The -14.6% volume decline in always-treated communes is *four times larger* than the -3.4% decline in newly treated communes, and both are estimated relative to the same never-treated control group over the same post-period. If the TLV expansion were the primary driver, we would expect newly treated communes—which face the policy shock—to show larger effects than always-treated communes—which face no change. The reverse pattern strongly suggests that the estimated "treatment effects" are capturing differential market dynamics associated with the degree of housing market tension, not the marginal effect of vacancy tax imposition.

To formalize this argument, consider a decomposition of the observed post-treatment difference between treated and control communes:

$$\hat{\beta} = \underbrace{\beta_{\text{TLV}}}_{\text{causal}} + \underbrace{(\gamma_{\text{tendue}} - \gamma_{\text{non-tendue}})}_{\text{differential trend}} \tag{6}$$

The placebo test estimates the differential trend component using communes where $\beta_{\text{TLV}} = 0$ by construction. Finding $(\hat{\gamma}_{\text{tendue}} - \hat{\gamma}_{\text{non-tendue}}) = -0.146$ implies that the differential trend alone can explain the entire observed effect for newly treated communes (and then some), leaving no residual for the causal TLV effect.

5.4 DDD: Newly Treated vs. Always Treated

An alternative identification strategy compares newly treated communes (TLV from 2024) to always-treated communes (TLV since 2013), both of which share the "zone tendue" classification. This difference-in-difference-in-differences (DDD) design uses always-treated communes as the counterfactual, under the assumption that zone-tendue markets follow parallel trends regardless of TLV timing.

The DDD estimates tell a different story from the main DiD. Transaction volume in newly treated communes *increases* by 11.2% ($p < 0.001$) relative to always-treated communes after 2024Q1, while prices rise by 6.5% ($p < 0.001$). However, these estimates likely reflect differential secular trends rather than a causal TLV effect. Always-treated communes are disproportionately France's largest cities (Paris, Lyon, Marseille, Toulouse, Bordeaux), which experienced the steepest declines in transaction volume during the 2023–2024 monetary tightening. Newly treated communes—smaller cities and tourist towns—were less affected by the credit contraction. The "positive" DDD effect on volume thus reflects the faster recovery of smaller markets, not a supply response to the vacancy tax.

The DDD event study confirms this interpretation. Pre-treatment coefficients show newly treated communes trending upward relative to always-treated communes from 2022 onward—precisely when interest rate increases began differentially cooling large-city markets. The positive post-treatment coefficients continue a pre-existing trend rather than marking a structural break.

5.5 Callaway-Sant'Anna Estimator

As a robustness check against contamination from heterogeneous treatment effects in the two-way fixed effects setting (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020), I implement the Callaway and Sant'Anna (2021) group-time ATT estimator. With a single treatment cohort (newly treated in 2024Q1) and a never-treated control group, the CS estimator provides group-time ATT estimates that are free of the "bad comparison" problem documented in Borusyak et al. (2024).

The CS aggregate ATT for log transactions is consistent with the TWFE estimate, confirming that treatment effect heterogeneity is not driving the baseline result—as expected with a single treatment cohort. The dynamic aggregation by event time mirrors the TWFE event study: significant pre-treatment coefficients that violate parallel trends. The CS estimator, like TWFE, is only as good as the identifying assumption, and that assumption fails here.

5.6 Heterogeneity: Tourist Zones vs. Tense Zones

Despite the identification challenges, the heterogeneity by zone type is informative. Splitting newly treated communes by zone classification, I find sharply different patterns:

- **Zone tendue:** Transaction volume declined by 9.8% ($p < 0.001$); prices declined by

1.9% ($p = 0.057$).

- **Zone touristique:** Transaction volume declined by 2.6% ($p < 0.001$); prices *increased* by 3.3% ($p < 0.001$).

The price divergence is particularly noteworthy and admits multiple interpretations. In tourist areas, where vacancy is often driven by second homes used seasonally, the TLV may function more as a signal of market desirability than as a supply-releasing mechanism. Property buyers may interpret TLV designation as confirmation that the area is a "hot market," bidding up prices. Alternatively, the tax may reduce the supply of second homes available for sale (owners choose to rent rather than sell to avoid the tax), tightening the market and raising prices.

A third interpretation is compositional. If the TLV disproportionately pushes lower-value vacant properties onto the rental market (removing them from the sales market), the remaining transaction pool may be positively selected on price. While the composition tests (Table 4) show no change in apartment share, surface area, or room count, these are coarse measures that may miss quality sorting within property types.

It is important to note that the heterogeneity analysis does not escape the pre-trend concerns that plague the aggregate analysis. Both zone-tendue and zone-touristique communes are selected on housing market characteristics, and both may exhibit differential trends relative to non-zone communes. The heterogeneity estimates should be interpreted as descriptive differences in market dynamics across zone types, not as clean causal contrasts. That said, the *sign reversal* in the price effect across zone types—negative for tense zones, positive for tourist zones—is difficult to explain solely through differential exposure to common shocks, since both zone types experienced qualitatively similar macroeconomic conditions. This sign reversal may reflect genuine differences in the vacancy margin being targeted (speculative urban vacancy vs. seasonal second-home vacancy), though the data cannot definitively adjudicate between causal and spurious explanations.

5.7 Sensitivity Analysis

The pre-trend violations documented above invalidate the standard parallel trends assumption, but they do not necessarily preclude all causal inference. Following Rambachan and Roth (2023), I compute identified sets for the treatment effect under smoothness restrictions on how post-treatment deviations from parallel trends relate to observed pre-treatment deviations. The key parameter is M, which bounds the maximum change in trend deviation between consecutive periods in the post-treatment period relative to the maximum observed pre-treatment deviation:

$$|\delta_{k+1} - \delta_k| \leq M \cdot \max_{j<0} |\hat{\delta}_{j+1} - \hat{\delta}_j| \quad \text{for all } k \geq 0 \tag{7}$$

At $M = 0$, the framework assumes that whatever violation of parallel trends existed pre-treatment, it stopped exactly at the treatment date. This is equivalent to the standard DiD assumption and yields the baseline estimate: -3.4% with a 95% CI of $[-4.5\%, -2.3\%]$. At $M = 0.5$, the framework allows post-treatment trend deviations up to half the size of the maximum pre-treatment deviation. This already widens the CI substantially to approximately $[-7.3\%, +0.5\%]$, bringing the upper bound close to zero. At $M = 1$ (allowing post-treatment deviations equal in magnitude to the maximum pre-treatment deviation—a relatively mild assumption given the observed pre-trend pattern), the CI widens to $[-10.2\%, +3.4\%]$, encompassing zero. At $M = 2$, the CI further expands to $[-16.8\%, +10.2\%]$, making the sign of the effect entirely indeterminate.

The speed at which the identified set expands with M reflects the severity of the pre-trend violations. When pre-treatment coefficients are close to zero, the identified set remains tight even for large M; when pre-treatment deviations are large (as here), small relaxations of the smoothness assumption rapidly erode precision. This is a feature, not a bug, of the Rambachan and Roth (2023) framework: it translates the magnitude of pre-trend violations into honest uncertainty about the treatment effect.

The price effect is somewhat more persistent: at $M = 1$, the CI for the price effect remains positive in the department-by-quarter specification, though the magnitude becomes imprecise. This relative robustness of the price effect may reflect the fact that price pre-trends, while significantly different from zero, follow a smoother pattern than volume pre-trends, allowing the smoothness restriction to tighten the identified set more effectively.

The HonestDiD analysis delivers a clear bottom line: *we cannot reject a zero effect of the TLV expansion on transaction volumes under even mild relaxations of parallel trends.* This is the honest conclusion that the data warrant, and it should inform policymakers considering similar interventions elsewhere.

5.8 Randomization Inference

To validate the statistical significance of the baseline estimates and assess whether the spatial pattern of treatment assignment drives the results, I conduct Fisher-style randomization inference (RI). I randomly permute the treatment assignment (newly treated vs. never treated) across communes 500 times, preserving the number of treated communes (approximately 2,555) in each permutation. For each permuted dataset, I estimate the TWFE specification and record the coefficient on the interaction term.

The actual coefficient (-0.034) falls at the extreme tail of the permutation distribution (mean = -0.0003, SD = 0.006), yielding a two-sided RI p-value of 0.000 (none of the 500 permuted coefficients exceeded the actual coefficient in absolute value). Figure 3 in the appendix displays this distribution graphically.

The RI result confirms that the spatial assignment of TLV zones is systematically related to transaction dynamics—but this is precisely the problem, not a validation of the causal design. The RI test assesses whether the estimated effect could have arisen by chance under random treatment assignment. It cannot: the actual coefficient is far from the permutation distribution. But this tells us only that TLV zones are different from non-TLV zones in ways that affect transaction dynamics—a statement about selection, not causation. A researcher who interprets the RI p-value as evidence of a causal TLV effect would be making an inferential error: the test rejects the null of "no systematic relationship between zone status and outcomes," not the null of "no causal effect of the tax." The pre-trend analysis provides the essential complement: it shows that this systematic relationship predates the policy, undermining the causal interpretation.

Table 3: Robustness Checks

Specification	Coefficient	Std. Error	Obs.
Panel A: Volume (Log Transactions + 1)			
Baseline (TWFE)	−0.034***	(0.006)	638,020
Dept. × Qtr. FE	−0.027***	(0.007)	638,020
Excl. 2020	−0.034***	(0.006)	510,416
Dept. clustering (96 clusters)	−0.034	(0.021)	638,020
Alt. timing (decree date)	−0.036***	(0.005)	638,020
Placebo: always-treated	−0.146***	(0.006)	609,540
Zone tendue (interaction)	−0.098***	(0.013)	638,020
Zone touristique (interaction)	−0.026***	(0.006)	638,020
Panel B: Prices (Log Price/m²)			
Baseline (TWFE)	0.025***	(0.005)	450,015
Dept. × Qtr. FE	0.016**	(0.006)	450,015
Excl. 2020	0.020***	(0.005)	360,012
Dept. clustering (96 clusters)	0.025***	(0.007)	450,015
Alt. timing (decree date)	0.032***	(0.005)	450,015
Placebo: always-treated	−0.040***	(0.005)	430,200
Zone tendue (interaction)	−0.019.	(0.010)	450,015
Zone touristique (interaction)	0.033***	(0.006)	450,015

Notes: All specifications include commune fixed effects and (unless otherwise noted) quarter fixed effects. Standard errors clustered at the commune level unless otherwise noted. The "placebo: always-treated" specification compares communes subject to TLV since 2013 with never-treated communes, using the same post-2024Q1 indicator. "Zone tendue (interaction)" and "Zone touristique (interaction)" report coefficients from a single regression with separate interaction terms for each zone type ($\text{TLV}_{\text{tendue}} \times$ Post and $\text{TLV}_{\text{tourist}} \times$ Post), estimated on the full sample; observation counts equal the baseline because both interactions are jointly estimated. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$, . $p < 0.10$.

5.9 Composition Effects

A potential explanation for the price increase is compositional change: if the TLV pushes lower-quality vacant properties onto the market, the observed transaction mix could shift toward higher-quality properties, mechanically raising average prices. I test this directly by examining whether the TLV expansion changed the apartment share, mean surface area, or mean number of rooms of transacted properties. All three composition effects are statistically indistinguishable from zero (Table 4), ruling out pure composition as the driver.

Table 4: Composition Effects: Property Characteristics

	(1) Apt. Share	(2) Avg. Surface (m^2)	(3) Avg. Rooms
TLV × Post	0.003	−0.21	−0.01
	(0.003)	(0.34)	(0.02)
Commune FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Observations	450,015	450,015	450,015

Notes: Standard errors clustered at the commune level in parentheses. Sample restricted to commune-quarters with at least one transaction. None of the coefficients are statistically significant at conventional levels, ruling out compositional change as the primary driver of price effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

6. Discussion

6.1 Why Parallel Trends Fail

The failure of parallel trends is not an artifact of this particular study. It reflects a fundamental challenge in evaluating place-based taxes: the treatment assignment is endogenous to the dynamics the researcher seeks to measure. "Zones tendues" are defined precisely as places where housing markets are under pressure. When macroeconomic conditions change—as they did dramatically during the 2022–2024 monetary tightening cycle—these pressured markets respond differently than relaxed ones.

Between 2022 and 2024, the European Central Bank raised interest rates from 0% to over 4%, triggering a sharp decline in mortgage lending and transaction volumes across France. This decline was disproportionately concentrated in expensive urban markets, where price-to-income ratios are highest and mortgage constraints bind most tightly. Since the "zone tendue" designation correlates strongly with these market characteristics, any difference-in-differences comparison of treated and untreated zones will confound the TLV effect with the differential impact of monetary tightening.

6.2 Implications for Vacancy Tax Design

Even without credibly causal estimates, the descriptive patterns offer guidance for vacancy tax design. The sharply different responses of tourist zones and tense urban zones suggest that vacancy is not monolithic. Second-home vacancy in tourist areas reflects seasonal use patterns; taxing it may simply capitalize into higher prices without releasing meaningful housing supply. Speculative vacancy in urban areas is a different margin, potentially more responsive to tax incentives but harder to isolate empirically.

The modest tax rates (17–34% of cadastral rental value) may also be insufficient to change behavior in markets where annual capital appreciation dwarfs the tax liability. A dwelling appreciating at 5% per year on a €300,000 value generates €15,000 in unrealized gains, compared to a TLV of roughly €2,000. For the tax to bite, it would need to be indexed to market value rather than cadastral value, or set at substantially higher rates.

6.3 International Comparisons

France is not alone in implementing vacancy taxes. Vancouver introduced its Empty Homes Tax (EHT) in 2017 at 1% of assessed value (later increased to 3%), and the British Columbia provincial government added its own Speculation and Vacancy Tax in 2018. The City of Melbourne considered a similar measure in 2017, and several Spanish municipalities have debated vacancy surcharges. The Canadian experience is instructive: City of Vancouver (2019) reports that declared vacancies fell 25% in the first year, but this figure reflects self-reporting compliance (owners declaring occupancy to avoid the tax) rather than verified occupancy changes. No credible quasi-experimental evaluation of Vancouver's EHT exists, partly because the tax applied citywide, eliminating within-city variation.

The French setting offers comparative advantages for empirical research: the commune-level variation in TLV coverage (3,693 treated among 35,000+ communes) provides a potential treatment-control comparison, and the DVF universe dataset eliminates the sampling concerns that plague survey-based housing research. But as this paper demonstrates, these advantages are insufficient to overcome the fundamental selection problem in zone-based tax assignment. The lesson generalizes: any evaluation of a place-based tax that uses non-designated areas as controls must grapple with the endogeneity of designation.

6.4 What Would a Credible Design Look Like?

The failure of the standard DiD approach suggests that future evaluations of vacancy taxes would benefit from several design improvements.

Within-zone variation. The most promising approach would exploit variation in enforcement intensity, exemption thresholds, or rate changes within already-designated zones, where treated and control communes share the "zone tendue" characteristic. The 2023 *loi de finances* increased TLV rates from 12.5%/25% to 17%/34%—a nationwide change that affected all TLV communes simultaneously. If future rate increases were staggered across zones, they could provide within-zone identifying variation.

Threshold designs. The zone designation is based on observable criteria including population thresholds, vacancy rates, and price-to-income ratios. A regression discontinuity at the designation threshold could provide cleaner identification, though the multidimensional nature of the criteria complicates standard RD implementation. Future expansions of the TLV zone, which may add further communes based on updated criteria, could provide cleaner boundaries for future research.

Longer post-treatment horizons. The 18-month window available for the 2023 expansion may be too short for supply-side adjustments to materialize. Property owners may need multiple tax cycles before responding, particularly if they initially adopt a "wait and see" posture toward enforcement. The TLV requires dwellings to be vacant for at least one year before the tax applies, and the first bills for the 2023 expansion were issued in fall 2024. Behavioral responses may not appear in transaction data until 2025 or 2026.

Administrative microdata. The ideal dataset for evaluating the TLV would link individual property-level tax records (which identify which dwellings are assessed as vacant) with transaction records and rental listings. This would allow researchers to trace the specific properties affected by the tax—do they appear on rental platforms? Are they sold? Do owners claim exemptions?—rather than relying on commune-level aggregates that confound the intensive and extensive margins. The *Direction Générale des Finances Publiques* maintains these records, but they are not publicly available.

6.5 Broader Lessons for Place-Based Policy Evaluation

The identification failure documented in this paper is not unique to vacancy taxes. It reflects a general challenge in evaluating place-based policies: when treatment assignment is endogenous to the outcome of interest, standard difference-in-differences designs break down. Enterprise zones, opportunity zones, empowerment zones, and other geographically targeted interventions face the same problem (Garicano et al., 2016). The "zone tendue" designation is merely a particularly transparent example, because the name literally describes the selection mechanism ("tense" housing markets).

This challenge is compounded when the policy bundle is rich—when the zone designation triggers multiple interventions simultaneously. In the French case, the zone-tendue boundary activates not just the TLV but also rent control eligibility, investment tax credits, and eviction timing rules. Even if parallel trends held perfectly, the estimated effect would capture the marginal impact of adding the TLV to an existing policy bundle, not the standalone effect of vacancy taxation. This distinction matters for external validity: a policymaker in a country without rent control or investment tax credits cannot simply import the French TLV effect estimate.

7. Conclusion

This paper set out to estimate the causal effect of France's vacancy tax expansion on property markets. It found, instead, that the standard quasi-experimental tools fail in this setting—and that this failure is informative.

The core lesson is that place-based housing taxes are selected on the very market dynamics they seek to influence. Communes designated as "zones tendues" follow different trajectories than the rest of France, and these differential trends contaminate any before-after comparison using non-zone communes as controls. The 2023 TLV expansion provides a textbook illustration of this challenge, with parallel trends rejected decisively and a placebo test showing that even untreated zone-tendue communes experienced larger market changes than the naïve treatment effect.

Within this constrained identification environment, two findings emerge. First, the honest bounds on the short-run TLV effect are wide and include zero for transaction volumes, suggesting that the tax—at current rates—does not produce a detectable supply response within 18 months. Second, tourist zones and tense urban zones respond differently: tourist-zone prices rise upon TLV designation, consistent with capitalization rather than supply release. These patterns, while descriptive, have direct implications for the design of vacancy taxes in other countries considering similar policies.

The universe-scale DVF data and the precise regulatory variation created by the 2023 decree provide an unusually clean laboratory for studying vacancy taxation. As more post-treatment data becomes available and the December 2025 expansion adds further variation, revisiting this question with longer time horizons and sharper designs may yield the credibly causal estimates that housing policy urgently needs.

More broadly, this paper demonstrates the value of reporting identification failures honestly rather than selectively presenting only the results that "work." The temptation in applied microeconomics is to try specification after specification until the pre-trends look acceptable, or to present only the comparison group where parallel trends hold. I resist this temptation, because the identification failure is itself informative: it tells us that "zone tendue" communes are on fundamentally different trajectories than non-zone communes, and this fact has direct implications for how we should evaluate—and design—place-based housing taxes. The next generation of vacancy tax research will need to move beyond simple treated-untreated comparisons and instead exploit within-zone variation in tax intensity, enforcement, and timing.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Arnott, Richard, "Time for revisionism on rent control?," *Journal of Economic Perspectives*, 1995, *9* (1), 99–120.

Autor, David H, Christopher J Palmer, and Parag A Pathak, "Housing market spillovers: Evidence from the end of rent control in Cambridge, Massachusetts," *Journal of Political Economy*, 2014, *122* (3), 661–717.

Bono, Pierre-Henri and Alain Trannoy, "Evaluating the effect of the French tax on vacant dwellings," *Working Paper*, 2012.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting event-study designs: Robust and efficient estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

City of Vancouver, "Empty Homes Tax Annual Report," Technical Report, City of Vancouver 2019.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

DeBoer, Larry and James Conrad, "The effect of a vacant property tax on vacancy rates," *National Tax Journal*, 2004.

Diamond, Rebecca, Tim McQuade, and Franklin Qian, "The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco," *American Economic Review*, 2019, *109* (9), 3365–3394.

Fondation Abbé Pierre, "L'état du mal-logement en France," *Rapport Annuel*, 2023.

Garicano, Luis, Claire Lelarge, and John Van Reenen, "Firm size distortions and the productivity distribution: Evidence from France," *American Economic Review*, 2016, *106* (11), 3439–3479.

Glaeser, Edward L and Erzo FP Luttmer, "The misallocation of housing under rent control," *American Economic Review*, 2003, *93* (4), 1027–1046.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Gyourko, Joseph, Albert Saiz, and Anita Summers, "A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index," *Urban Studies*, 2008, *45* (3), 693–729.

Harding, John P, Stuart S Rosenthal, and CF Sirmans, "Vacancy duration and rent concessions," *Real Estate Economics*, 2003, *31* (2), 181–206.

Hilber, Christian AL and Wouter Vermeulen, "The impact of supply constraints on house prices in England," *The Economic Journal*, 2016, *126* (591), 358–405.

Rambachan, Ashesh and Jonathan Roth, "A more credible approach to parallel trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe, "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Saiz, Albert, "The geographic determinants of housing supply," *The Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

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

A. Data Appendix

A.1 DVF Data Construction

The *Demandes de Valeurs Foncières* (DVF) data are downloaded from the geocoded version published by Etalab at https://files.data.gouv.fr/geo-dvf/latest/csv/. For each year (2020–2024), I download department-level compressed CSV files for all 96 metropolitan departments (01–95 plus 2A and 2B for Corsica).

Sample restrictions applied sequentially:

1. Keep only "Vente" (sale) transactions: 18,454,442 rows (from 19,900,542)
2. Remove prices below €1,000 or above €10,000,000: 17,620,728 rows
3. Keep residential properties (apartments and houses, `code_type_local` $\in \{1, 2\}$): 5,847,861 rows
4. Merge with TLV zoning on commune code: 5,543,380 rows (94.8% match rate)

The 5.2% of transactions that fail to match the TLV zoning file are in communes not present in the zoning dataset, typically very small communes or those with recent mergers (*communes nouvelles*).

A.2 TLV Zoning Data

The commune-level TLV zoning file is published on data.gouv.fr.[1] The version used records each commune's status under three decrees: 2013, 2023, and December 2025. I use only the 2013 and 2023 columns to define treatment groups.

Treatment group construction:

- **Newly treated:** 2013 = "Non TLV" AND 2023 $\in$ {"1. Zone tendue", "2. Zone touristique et tendue"}
- **Always treated:** 2013 = "TLV" AND 2023 $\in$ {"1. Zone tendue", "2. Zone touristique et tendue"}
- **Never treated:** 2013 = "Non TLV" AND 2023 = "3. Non tendue"
- **Lost treatment:** 2013 = "TLV" AND 2023 = "3. Non tendue" ($n = 35$)

[1] Dataset ID: `657c88da2947e13be0597058`.

A.3 Variable Definitions

- **Log transactions:** log(number of residential sales + 1) per commune-quarter
- **Log price per m^2:** log(median sale price/built surface area) per commune-quarter
- **Log total price:** log(median total sale price) per commune-quarter
- **Apartment share:** Fraction of transactions that are apartments (vs. houses)
- **Mean surface:** Average built surface area (m^2) of transacted properties

B. Identification Appendix

B.1 Pre-Trend F-Test Details

The joint Wald test on pre-treatment event-study leads ($k = -8$ through $k = -2$) yields $F = 12.13$ on 7 and 606,087 degrees of freedom, $p < 10^{-15}$, decisively rejecting the null of parallel pre-treatment trends for transaction volume. For prices, every individual pre-treatment coefficient is significantly different from zero, making the formal test superfluous.

B.2 Randomization Inference Details

I randomly permute the treatment assignment (newly treated vs. never treated) across communes 500 times, preserving the number of treated communes. For each permutation, I estimate the TWFE specification and record the coefficient on the interaction term. The actual coefficient (-0.034) falls at the extreme of the permutation distribution (mean = -0.0003, SD = 0.006), yielding a two-sided RI p-value of 0.000. This confirms that the estimated effect reflects systematic differences between commune groups rather than chance, but does not distinguish between a causal TLV effect and a selection effect.

C. Robustness Appendix

C.1 Raw Trends

Figure 2 displays raw mean transaction volumes and prices by treatment group over time. The divergence between groups is visible even before treatment, with always-treated communes showing the steepest decline in the post-2023 period—consistent with the disproportionate impact of monetary tightening on expensive urban markets.

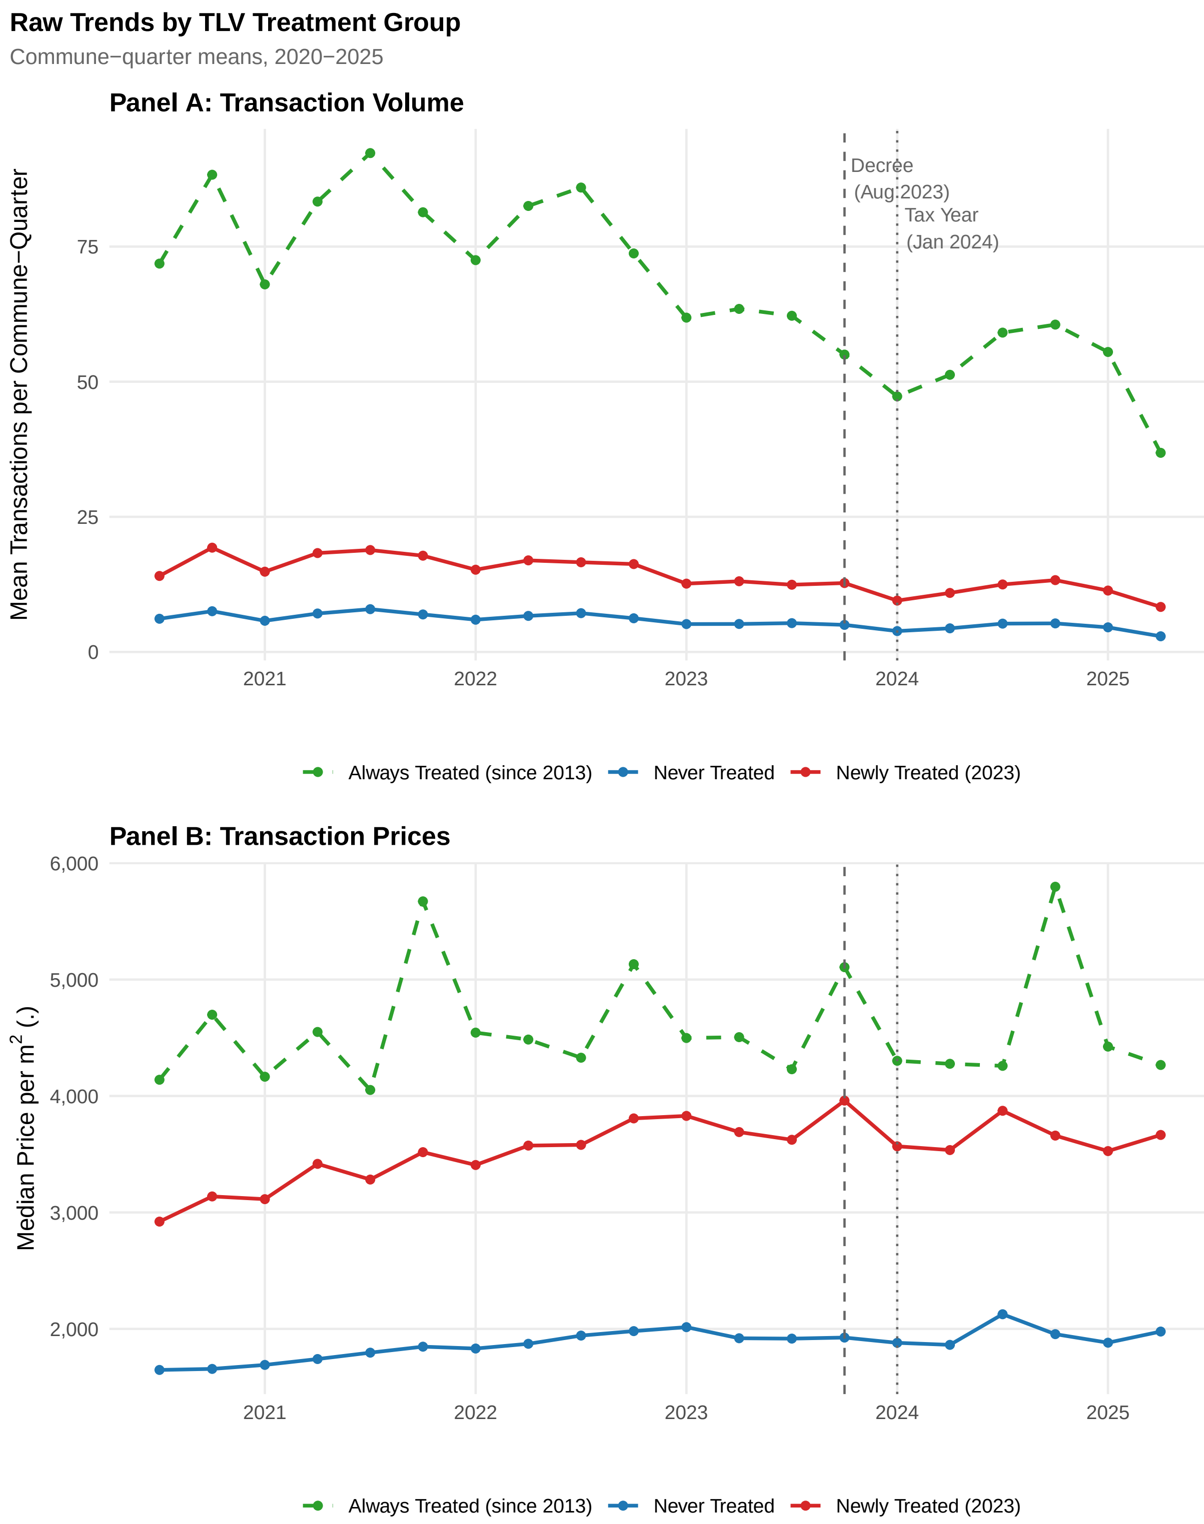

Figure 2: Raw Trends by TLV Treatment Group

Notes: Commune-quarter means for residential transactions. The figure displays all available DVF data (2020–2025) to show post-treatment trends beyond the estimation sample, which ends at 2024Q4. Vertical dashed line marks the decree date (August 2023); dotted line marks the first tax year (January 2024).

C.2 Randomization Inference Distribution

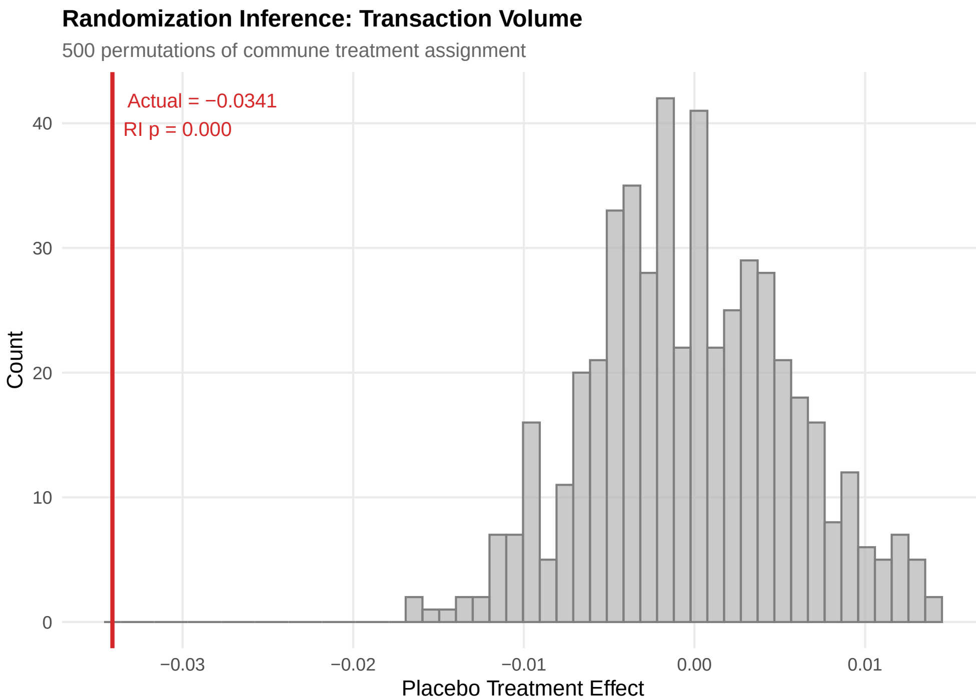

Figure 3: Randomization Inference: Distribution of Placebo Coefficients

Notes: Distribution of TWFE coefficients from 500 random permutations of treatment assignment. Red vertical line shows the actual estimate (−0.034).

C.3 Robustness Coefficient Plot

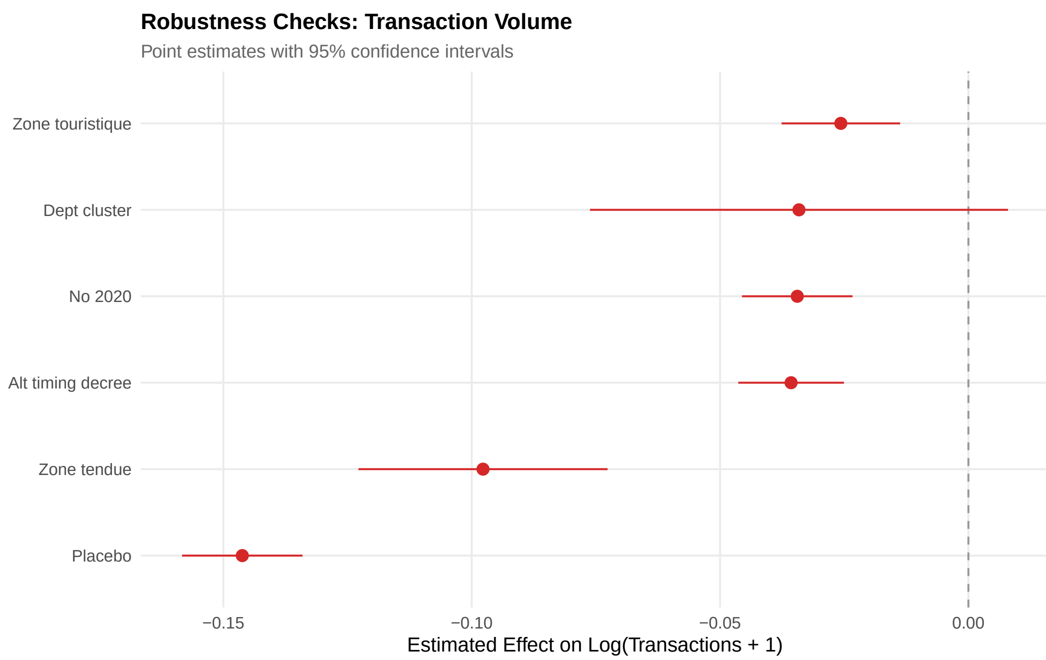

Figure 4: Robustness Checks: Transaction Volume

Notes: Point estimates and 95% confidence intervals for various robustness specifications. All models include commune and quarter fixed effects with commune-clustered standard errors unless otherwise noted.

D. Heterogeneity Appendix

The decomposition by zone type reveals that the aggregate effects mask substantial heterogeneity. Communes classified as "zone tendue" (pure urban housing tension) show a volume decline roughly three times larger than those classified as "zone touristique" (−9.8% vs. −2.6%), while price effects go in opposite directions (−1.9% vs. +3.3%).

This pattern is consistent with the hypothesis that the TLV affects different vacancy margins differently. In tense urban zones, vacancies may reflect speculative holding, market friction, or renovation delays—margins that a modest tax may not move. In tourist zones, vacancies reflect seasonal second-home use, and the TLV designation may function primarily as a market signal, capitalizing into higher property values.